



**How to design, optimise and implement a
fibre-tip Fabry-Pérot cavity for quantum
networks of atoms and photons**

A thesis submitted for the degree of
Doctor of Philosophy (DPhil)

Atomic & Laser Physics
St Hugh's College, University of Oxford

Marwan Mohammed

Hilary term 2021

How to design, optimise and implement a fibre-tip Fabry-Pérot cavity for quantum networks of atoms and photons

Marwan Mohammed

A thesis submitted for the degree of
Doctor of Philosophy (DPhil), Atomic & Laser Physics
St Hugh's College, University of Oxford

One of the most exciting prospects of quantum physics today is to simulate large quantum systems, intractable by classical computation, by connecting many smaller systems together. A cavity-based quantum network of atoms and photons constitutes one such approach, where each cavity interfaces between stationary atomic qubits in a node and photonic qubits travelling along optical fibre channels. For this scheme to work well, the cavity-based interface must be faithful and well-controlled.

In this thesis, we construct a fibre-tip Fabry-Pérot cavity by aligning two optical fibre-tips against each other, where each has a concave, mirror-coated surface. Three properties make it highly suitable for interfacing: a small mode volume for strong atom-photon coupling, directly fibre-coupled cavity light, and optical access for trapped atoms. For the latter, we design a scheme that allows for a 2D reconfigurable array of atoms, held in the cavity mode by holographically-generated optical tweezers. Our particular design will allow us to move atoms independently of one another.

We find the atom-cavity system parameters for optimal single photon generation, absorption and remote state transfer using V-STIRAP (vacuum-stimulated emission by Raman adiabatic passage). Using master equations and the Nelder-Mead algorithm, we optimise both the parameters and V-STIRAP pulse shapes to do this.

We assemble the first fibre cavity formed of two single-mode fibres in ultra-high vacuum. Aligning two single-mode fibres is particularly demanding; we overcome this with an assembly method that gives permanent sub-micron alignment precision, unlike other similar cavities that need realigning using a translation stage. We fully characterise the cavity's parameters, then test it by probabilistically loading atoms from a magneto-optical trap, measuring the Purcell-enhanced fluorescence from the cavity.

Altogether, our advances facilitate deterministically controlled, strongly-coupled and networked quantum systems.

Statement of Originality

I hereby declare that this thesis is entirely my own work.

Marwan Mohammed
Hilary term 2021

Contents

Abstract	ii
Contents	iv
1 Introduction	1
1.1 Structure of this thesis	5
2 Theory of cavity quantum electrodynamics	6
2.1 Optical cavity	8
2.1.1 Fabry-Pérot resonator	9
2.1.2 Resonator modes	11
2.2 Atom-cavity interactions	17
2.2.1 Atom-cavity coupling	18
2.2.2 Losses: damped atom-cavity system	29
2.2.3 Single photon source	35
2.3 Summary	42
3 Optimal parameters for cavity-based quantum networks of atoms and photons	43
3.1 Introduction	44
3.2 Simulation elements	47
3.3 Optimal atom-cavity parameters	50
3.4 Krotov pulse shape optimisation	57
3.4.1 Setting up the Krotov algorithm	58
3.4.2 Performing pulse shape optimisation with Krotov	62
3.5 Nelder-Mead pulse shape optimisation	65
3.6 Optimising remote state transfer	73
3.7 Conclusion	78
4 Making a Fibre-tip Fabry-Pérot cavity for fully deterministic quantum interfacing	81
4.1 How are fibre-tip mirrors made?	83
4.2 Mode-matching	87
4.3 How accurately can we specify a fibre-tip mirror?	96
4.3.1 Decentration	96

4.3.2	Ellipticity/Birefringence	98
4.3.3	Clipping loss	100
4.4	Assembling a fibre Fabry-Pérot cavity	103
4.4.1	Cavity alignment/characterisation apparatus	103
4.4.2	Cavity testing before vacuum	108
4.4.3	Cavity mount assembly	116
4.4.4	Counteracting misalignment caused by epoxy deformation	122
4.5	Putting the cavity into vacuum	126
4.5.1	First attempt	127
4.5.2	Second attempt	129
4.6	Conclusion	135
5	Purcell-enhanced few-atom fluorescence in a fibre cavity	137
5.1	Experimental map and overview	138
5.2	Experimental design and components	140
5.2.1	Optical tweezers	140
5.2.2	Magneto-optical trap	142
5.2.3	Vacuum chamber	153
5.2.4	Lasers and locking	159
5.3	Cavity stabilisation	167
5.3.1	Mechanical stability of the cavity	167
5.3.2	Frequency stabilisation of the cavity	173
5.4	Experiment methods introduction	176
5.5	Beam of atoms: pushing atoms into the cavity	177
5.5.1	Results of atom-pushing experiment	177
5.6	Probabilistic loading: dropping atoms through the cavity	179
5.6.1	Loading the cavity with atoms	179
5.6.2	Driving beam	183
5.6.3	Atom loading sequence	184
5.6.4	Results of probabilistic loading	185
5.7	Conclusion	189
6	Perspectives on this work	191
6.1	Other implementations of quantum networks	192
6.2	Present and future fibre cavity research	195
7	Conclusion	201
A	How to calculate the mixing angles of two- and three-level systems	206
A.1	Two-level system	206
A.2	Three-level Lambda-system	208

B	Liouville space	212
C	Theory of a 2D array of single-atom optical tweezers	216
C.1	Theory of single-atom optical tweezers	216
C.2	Generating a reconfigurable 2D array of single atoms	220
	Bibliography	223

List of Figures

1.1	Quantum network of fibre cavities.	4
2.1	A cavity resonator and a cricket ground.	7
2.2	Fabry-Pérot cavity sketch and transmission properties.	10
2.3	Shape of a fundamental resonator mode.	13
2.4	Energy levels of a 2-level atom and a separate cavity.	19
2.5	Jaynes-Cummings ladder for an atom-cavity system, resonant case.	27
2.6	Jaynes-Cummings spectrum for an atom-cavity system, off-resonant case.	30
2.7	Loss processes in an atom-cavity system.	30
2.8	2-level atom and cavity dynamics in bad cavity and strong coupling regimes.	34
2.9	3-level atom-cavity Lambda-system.	36
2.10	3-level atom-cavity system dynamics for single photon emission.	41
3.1	Using the Nelder-Mead algorithm to optimise the atom-cavity system parameters for photon generation.	53
3.2	Convergence of atom-cavity parameter optimisation.	55
3.3	Scan of detunings from optimal atom-cavity parameters.	56
3.4	Simulated photon absorptions for an atom-cavity system.	57
3.5	Pictorial representation of Krotov’s algorithm.	63
3.6	Pulse shape updates and system dynamics for one iteration of Krotov.	64
3.7	How we use coherence and splines to optimise pulse shapes using Nelder-Mead.	68
3.8	Optimal laser pulse shapes for single photon emission and absorption.	70
3.9	Optimal pulses for emitting photons efficiently within a time window.	72
3.10	Optimal laser pulse shapes for remote state transfer between two cavities.	75
3.11	Remote state transfer optimisation with transmission/loss channels modelled.	77
3.12	Remote state transfer optimisation in a hybrid cavity network.	78
4.1	Envisaged fibre cavity quantum interface.	82
4.2	Construction stages of a fibre Fabry-Pérot cavity.	84
4.3	Illustration of the mode mismatch between fibre and cavity.	87
4.4	Illustration of the modes involved in our mode-matching theory.	88
4.5	Mode-matching efficiency versus cavity length and misalignments.	94
4.6	Illustrations of ideal and non-ideal fibre-tip ablations.	96
4.7	Fibre-tip topography using an atomic force microscope.	98

4.8	Birefringence of our fibre Fabry-Pérot cavity.	101
4.9	Cavity characterisation apparatus and fibre-tip alignment stage.	105
4.10	Fibre-tip mirror coating's transmission versus wavelength.	106
4.11	Example of a cavity resonance linewidth measurement.	107
4.12	Finesse versus length of many fibre-tips aligned against a nearly flat mirror.	110
4.13	Finesse versus length of the final cavity, and bend loss as the fibres are coiled.	112
4.14	Mode-matching and asymmetry measurements of the final cavity shortly before it is glued.	115
4.15	Cavity assembly method.	117
4.16	3D model of the cavity mount.	119
4.17	Microscope image of V-groove alignment.	120
4.18	Assembled fibre Fabry-Pérot cavity.	121
4.19	Epoxy curing mechanism and bakeout simulation tests.	124
4.20	Microscope images of a fibre cavity that failed in vacuum.	129
4.21	Thermal self-locking in an earlier fibre cavity in vacuum.	130
4.22	Photo of V-groove taper.	131
4.23	Scanning a free-spectral range for the assembled fibre cavity.	133
4.24	Assembled fibre cavity in vacuum cell.	135
5.1	Map of the experimental setup used for measuring atom-cavity interactions.	139
5.2	Laser beam arrangement for the cavity in the vacuum cell.	141
5.3	Laser cooling of atoms: Doppler cooling mechanism and optical molasses.	144
5.4	Magneto-optical trap, and Zeeman shift from coils.	148
5.5	Photo of coils and beams for the magneto-optical trap.	150
5.6	Photo of vacuum apparatus.	154
5.7	Cavity mounting finger and holding assembly for vacuum chamber.	158
5.8	Hyperfine level structure of the ^{87}Rb D_2 line, and what transitions we use.	161
5.9	Error signal for frequency locking to an atom, and electronics used to get it.	165
5.10	Electronics for beat-note offset lock.	167
5.11	Power spectrum of the cavity transmission in vacuum without scanning.	170
5.12	Mechanical vibration response of the cavity in vacuum.	171
5.13	Simulated mechanical vibration response of the cavity mounting finger.	173
5.14	Histograms characterising how we mitigate the effect of vibrations when we activate the feedback from a lock-in amplifier.	175
5.15	Polarisation drift of the cavity-collected fluorescence from an atomic beam.	178
5.16	Images of the magneto-optical trap for our cavity loading scheme.	180
5.17	Timings of the cavity loading sequence.	185
5.18	Histograms of fluorescence counts showing cavity Purcell enhancement.	187
6.1	Possible network topologies for a quantum network of many fibre cavities.	198
C.1	How to holographically generate a 2D array of trapped atoms.	221

List of Tables

3.1	Initial and optimised values, and bounds, for the atom-cavity parameters when using the Nelder-Mead algorithm.	52
4.1	Specification of the fibre-tip mirrors.	85
4.2	Comparison of epoxies subjected to simulated bakeouts.	126
4.3	Summary of cavity parameters.	134

Symbols

In general, we denote vectors \mathbf{a} , quantum states $|a\rangle$, and quantum operators \hat{a} , and we define symbols upon first use in a chapter. This list is not comprehensive: we list only frequently used symbols and/or those which could be misinterpreted (e.g. certain subscript variations).

\mathbf{B}	magnetic field vector	I	electric field intensity
c_0	speed of light in vacuum	I_{sat}	saturation intensity
C	co-operativity	$\hat{\mathbb{1}}$	identity operator
\mathbb{C}	cost function	J	performance function
\hat{c}, \hat{c}^\dagger	cavity photon annihilation, creation operator	J_T	final-time performance function
\hat{C}_n	collapse operator	k	wavevector
d	decentration (length)	L_{cav}	cavity length
D	atomic dipole moment; diameter	\mathcal{L}	mirror loss intensity
\hat{D}	dissipation superoperator	$\hat{\mathcal{L}}$	Liouvillian superoperator
e	charge of an electron	M_F	z -component of hyperfine angular momentum
$ e\rangle$	atom's excited state	\hat{n}	cavity photon number operator
\mathbf{E}	electric field vector	q	complex distance from beam waist
\mathbf{E}_0	amplitude of electric field	r	mirror reflection amplitude; radial distance
E_{00}	electric field of fundamental transverse mode	R	mirror reflection intensity
F	fidelity; hyperfine total angular momentum; transfer function	\mathbf{r}	position vector
F_P	Purcell factor	\hat{R}	reflection operator (classical)
F_x, F_z	force acting in 1D	S	laser pulse shape
\mathcal{F}	finesse	t	mirror transmission amplitude; time
g	atom-cavity coupling strength; acceleration due to gravity	T	mirror transmission intensity; temperature
$ g\rangle$	atom's ground state	T_d	dipole trap depth (temperature)
h	Planck constant; height	\hat{U}	evolution operator
\hbar	reduced Planck constant	\hat{U}_0	evolution operator in interaction picture
\hat{H}	Hamiltonian/energy operator	U_d	dipole trap potential
\hat{H}_{int}	Hamiltonian of an interaction	$ u\rangle$	atom's ground state
\hat{H}_I	Hamiltonian in interaction picture		

V_{eff}	effective mode volume	σ	MOT width ($1/e^2$)
V_m	mode volume	σ^\pm	label for atomic transition where $\Delta M_F = \pm 1$
w	beam radius		
w_0	beam waist	$\hat{\sigma}^+, \hat{\sigma}^-$	atomic raising, lowering operator
w_f	fibre's mode waist	$\hat{\sigma}_Z$	Pauli-Z operator
w_m	mode (beam) radius at mirror	Υ	spontaneous emission number
z	position along cavity (or quantisation) axis	ϕ^\pm	cavity mode fields (classical)
		$\hat{\phi}_{\text{in}}$	photon in-coupling operator
z_R	Rayleigh length	Φ	mixing angle of three-level system
α	complex overlap amplitude (mode-matched part)	ψ^\pm, ψ_r	fibre mode fields (classical)
α_G	Gouy phase	$ \phi^\pm\rangle$	classical cavity mode integrals
β	coherence; complex overlap amplitude (fibre-interfering part)	$ \psi^0\rangle$	three-level system's dark eigenstate
γ	atomic amplitude decay rate	$ \psi^\pm\rangle$	three-level system's bright eigenstates; classical fibre mode integrals
Γ	atomic intensity decay rate		
Δ_B	cavity's birefringent splitting	ω	angular frequency
Δ_C, Δ_L	cavity, laser detuning	ω_a	atomic transition frequency
$\Delta\omega_{\text{FSR}}$	free spectral range	ω_c	cavity resonance frequency
$\Delta\omega_{\text{FWHM}}$	cavity linewidth (FWHM)	ω_L	laser frequency
ϵ	mode-matching efficiency	Ω_0	peak laser Rabi frequency
$\hat{\epsilon}$	polarisation direction unit vector	Ω_L	laser Rabi frequency
ϵ_0	electric constant		
ζ	radius of curvature	$ +\rangle, -\rangle$	two-level system's dressed states
η	photon generation/absorption efficiency; number density of atoms	∇	vector differential operator
θ	mixing angle of two-level system; geometric angle		
Θ	mixing angle of three-level system		
Θ_{div}	divergence		
κ	cavity field decay rate		
λ	wavelength; inverse step size		
μ_0	magnetic constant		
π	constant; label for atomic transition where $\Delta M_F = 0$		
Π	window function		
ρ	radial distance from cavity axis; density		
$\hat{\rho}$	density matrix		
$ \hat{\rho}\rangle\rangle$	state in Liouville space		
ϱ	density of photon states		

Acronyms

AC alternating current	LED light-emitting diode
AFM atomic force microscope	MM multi-mode
AOM acousto-optic modulator	MOT magneto-optical trap
AR anti-reflection	MTS modulation transfer spectroscopy
AWG arbitrary waveform generator	
CQED cavity quantum electrodynamics	NA numerical aperture
CRAB chopped random basis	ND neutral-density
CW constant power	NPBS non-polarising beam-splitter
DC direct current	
ECDL external cavity diode laser	PBS polarising beam-splitter
EIT electromagnetically induced transparency	PDH Pound-Drever-Hall
EM electromagnetic	PID proportional-integral-derivative
EOM electro-optic modulator	PM polarisation-maintaining
	QWP quarter-wave plate
FEM finite element analysis	RF radio frequency
FFPC fibre Fabry-Pérot cavity	RMS root-mean-square
FIB focused ion beam	RWA rotating wave approximation
FM frequency modulation	
FORT far-off-resonant trap	SLM spatial light modulator
FPC Fabry-Pérot cavity	SM single-mode
FPI Fabry-Pérot interferometer	SPCM single-photon counting module
FSR free-spectral range	STIRAP stimulated emission by Raman adiabatic passage
FWHM full-width at half-maximum	
FWM four-wave mixing	UHV ultrahigh vacuum
GRAPE gradient ascent pulse engineering	UV ultraviolet
HVA high-voltage amplifier	V-STIRAP vacuum-STIRAP
HWP half-wave plate	
IBS ion beam sputtering	WLI white light interferometry
IR infrared	ZPE zero-point energy

Glossary

coherence A well-defined phase relationship between the constituent states of a superposition state. A pure quantum state—one fully described by a single superposition state only—can be said to be fully coherent (note that a state with a single constituent is also a pure state).

decoherence The loss of coherence of a state; the process through which a pure (quantum) state becomes a (classical) mixture of states. An irreversible interaction of a system with its environment occurs. This may be thought of as perturbing the phases that define the coherent superposition with random kicks.

lock Fix a parameter to a well-defined reference value, usually in order to eliminate parameter fluctuations. *Example:* ‘We locked the laser frequency to the atomic resonance.’

scan Change a parameter back and forth across a range, usually across a central value of interest. *Example:* ‘We scanned the cavity frequency over the atomic resonance.’

superoperator An operator which maps a regular operator to another operator: $\hat{S}\hat{A} \mapsto \hat{B}$. Practically, this is simply an operator upon which we multiply on the left and right. *Example:* see equation 3.2. (Note: a regular operator maps a state vector to another state: $\hat{A}|a\rangle \mapsto |b\rangle$).

Introduction

Nature is quantum. For more than a century, quantum physics has proved to be phenomenally successful in correctly explaining our observations. Street lights are yellow because their sodium atoms only emit light at certain wavelengths. A photon behaves like a particle as it hits a detector, but like a wave as it interferes with itself in an interferometer. Because our quantum world is so fundamental, we wish to understand it better.

How do we do so? One thing we may want to do is simulate quantum systems. In 1982, Feynmann stated that a quantum simulator is more efficient than a classical one at simulating quantum systems [1]. To see why this is true, consider a state comprised of N quantum bits or qubits. A classical bit is either 0 or 1, so there are 2 possible measurement outcomes. For N bits, we can measure N outcomes at once by reading the value of each bit. However, a quantum bit can be in a superposition of 0 and 1. Therefore, if we measure an N -qubit state, we act on all the possible combinations of classical bits at once, comprising 2^N outcomes.

There are myriad real-world applications of such simulations or computations: for example, we can simulate complex chemistry at the quantum level to design more effective drugs and medicines [2], and exponentially speed up computations such as path-finding and factorisation [3], the latter of which may break current cryptographic codes [4]. Recently, a quantum computer solved a problem intractable to classical computers [5]: it achieved

‘quantum supremacy’ [6].¹ However, reaching this goal needed massive advances such as fault tolerance and near-lossless operation [8–12]. Quantum physics tells us why: qubits are sensitive to their environment, and interacting with it may erode their quantum properties. Systems made of a few qubits perform well [13], but in general, the larger the system, the harder it becomes to isolate qubits from their environment and perform logic on it faithfully.

Because small systems perform well, then instead of making one large system, we can connect many small systems together. As long as the connections use qubits, we can call this a quantum network [14]. Typically, we use stationary qubits at the network’s nodes—for example, an atom in a particular electronic state—and photonic qubits travelling along optical fibres between the nodes [15]. Stationary qubits can be well-isolated from their environment, and so we can store and process our quantum information accurately within the nodes. Meanwhile, photons are well-suited to quantum communication between the nodes, because their interactions with the environment and with other photons are weak. Therefore, they maintain the information well, although they do not stay in one place like the stationary qubits.

With larger systems, we can also create larger entangled states. Entanglement is a crucial resource in quantum networks. If the state of two or more qubits cannot be written as separate states of single qubits, they are entangled.² This implies that a measurement outcome of a single qubit may be correlated to another qubit: imagine flipping two separate coins, and knowing that if one turns up heads, the other must also turn up heads.³ At a high level, we need to create entanglement to simulate entangled systems. Furthermore, entanglement allows us to use quantum teleportation [16] and quantum repeaters [17] to send qubits faithfully

¹ Since we have only just reached the cusp of quantum supremacy, researchers are debating whether or not current classical systems could achieve the same performance [7].

² Mathematically, they are entangled if we cannot write the combined state as the tensor product of individual qubit states.

³ We have assumed we already know what the entangled state of the two coins is, so we do not gain any information with our measurement.

over long distances. Highly entangled states also allow for quantum-enhanced sensing [18], distributing secure cryptographic keys [19], quantum error correction [3], and also a quantum computing paradigm known as cluster state computing [20].

Altogether, a quantum network may contain the same information as one large quantum system, with distributed information, states and entanglement across the nodes, but ideally operating with greater fidelity and reliability because we can better isolate each node from its environment compared to a single large system.

Staying true to this vision, one approach is to build cavity-based quantum networks of atoms and photons [21]. Figure 1.1 illustrates such a network; the nodes and channels are represented by optical cavities and optical fibres respectively. Single atoms trapped in each cavity are the stationary qubits; we strike an atom through an open side of the cavity with a laser pulse, causing it to emit a photon. The photon then leaks out of the cavity, travels along the optical fibre, and into a second cavity, where another laser pulse causes that atom to absorb that photon. Here, the cavity is the key ingredient for our quantum interface [22]. Excitingly, elementary networks of two cavity-based nodes have already been well-established [23, 24]. In principle, this entire process can be coherent and reversible [25], and thus also the storage and delivery of qubits [26].

The cavity-based scheme is appealing because the strong coupling between the states of the atom and photon allow for all of the above to take place efficiently and faithfully [21]. For atom-photon interfacing, we may map between the atom's electronic state and the photon's polarisation state [23]: these are our logical qubits. Any computation or storage/delivery can be made resilient to photon loss by combining a heralding protocol with quantum repeaters [17, 27–29]: here, a non-destructive detection heralds the success of a state transfer. We can also create remote entanglement ‘by default’ after a successful remote state transfer [23], by




Figure 1.1: Quantum network of fibre cavities. **(a)**: Many-node network. One node consists of a trapped atom inside a cavity (inset), with optical tweezers—tightly focused laser beams that trap particles—illustrated as one method of trapping. Here, each node is connected to two others via two fibres. **(b)**: An elementary network, demonstrating how cavity-based quantum networks function. Two optical cavities connected by one fibre form an elementary network. A laser pulse strikes an atom, emitting a cavity photon. The photon leaks out, then travels via the fibre into another cavity, where another laser pulse causes another atom to absorb it. The entire process is coherent and reversible.

performing partial Bell measurements on pairs of photons coming from separate nodes [30],⁴ or by entangling photon pairs from different cavities with photonic chips and then measuring the photon pairs to transfer the entanglement to the atoms [31]. What remains as fundamental limitations on network performance are the coherence⁵ and efficiency of the processes [33]. So, we must minimise loss and decoherence.

Let us distinguish the effects of loss and decoherence. Loss causes a network operation not to happen. Decoherence means it happens, but with an undesired outcome. Losses result from mirror scattering or absorption,⁶ emission into free space, or fibre loss. Decoherence results

⁴ This has 50% maximum efficiency, because two of the four measurement outcomes cannot be distinguished from each other. Nevertheless, depending on the experimental parameters it might create entanglement more quickly.

⁵ For instance, coherence of photon(ic wave packet)s in space and time, or in polarisation [31, 32]. In a simpler way, coherence means we maintain the phase relationship between quantum states or their parts.

⁶ We will use stacks of dielectric coatings for the mirrors, as we will describe in chapter 4. Although the standing-wave cavity mode penetrates into the stacks, causing unavoidable losses [34], this only becomes apparent for cavities 2–3 orders of magnitude smaller than what we use in this thesis [35].

from spontaneous emission of the atom, incoherent laser control pulses, or imperfect atomic absorptions/emissions of the photon(s). Therefore, when we say that an ideal atom-photon interface must be efficient and coherent, these aspects above are what we want to minimise.

This brings us to the topic of this thesis: how do we design, implement and optimise a cavity-based quantum interface for quantum networks of atoms and photons? In the following chapters, we design a fibre Fabry-Pérot cavity (FFPC) for this purpose, where mirrors on the ends of two optical fibres form our cavity, trapping photons in a small volume. Each fibre-tip has a concave, mirror-coated surface, and we align the two against each other. Three properties make it highly suitable for interfacing: a small mode volume for strong atom-photon coupling, directly fibre-coupled cavity light, and optical access for trapped atoms. The latter property allows for perfect localisation of an atom in the cavity using lasers [36]. Thus, the goal of the thesis is to investigate how to design, implement and optimise our cavity for efficient, coherent atom-photon interfacing.

1.1 Structure of this thesis

Firstly, in chapter 2 we give the background to this thesis: we explain the theory of optical cavities with atoms, showing how the quantum interface works and what it is capable of. Next, in chapter 3 we run simulations to optimise cavity-based interfaces for quantum networking. In chapter 4, we design and build our FFPC quantum interface in the laboratory. Then in chapter 5, we present the full design of the rest of the experiment (including our design that integrates many atoms at once with the cavity). We then perform experiments demonstrating atom-cavity interactions and how suitable our FFPC is as a quantum interface. We give an outlook and perspective on our research in chapter 6, postulating future experiments and techniques that utilise our advances. Finally, we conclude in chapter 7.

Theory of cavity quantum electrodynamics

Suppose that a typical English-Summer deluge of rain floods The Oval cricket ground in London to a depth of 3 metres. Now suppose that an Australian spectator, entrepreneurial in spirit, rides a jet ski to the position just over the batting crease. Because the rules of cricket are complicated, England need to topple that wicket for some arbitrary reason: with the stumps shifting loosely in the ground, a slight tap from above will suffice. But between the flood, the jet ski and the Australian, what could England do? Sitting on the water's edge at the Pavilion, the team hatches a plan.

'On my signal, all of us splash our bats as hard as possible against the water,' orders England captain Joe Root. 'The waves will make the jet ski bob up and down in the water, and if we make large enough waves, the jet ski will bob down and knock the wicket from above.' Root rubs his hands together and grins like a poorly drawn Hanna-Barbera character.

'I'm not so sure,' interjects baby-faced all-rounder Craig Overton. 'By my calculations, I don't think our bats against the water will make large enough waves. We need to narrow the waves somehow to stop them spreading out away from the jet ski, and we need to time our splashes to match the natural frequency of the jet ski.' Overton hands him an envelope, on which his estimates are scribbled on the back. 'That way, our energy will be efficiently




Figure 2.1: **Left:** The Oval cricket ground, Kennington, London [37]. **Right:** An optical cavity resonator. Note that the ordering of left and right is intentional, reinforcing how difficult it can be to tell these two systems apart.

transferred. Then we might just stand a chance.’

‘You’re off the team, Craig.’

Notwithstanding the England team’s unorthodox method of play, their general idea is workable (although not sensible). The situation described is actually somewhat similar to the problems faced in cavity quantum electrodynamics: in this case, The Oval is a cavity resonator, the splashing waves correspond to the light field, and the jet ski is an atom. Figure 2.1 demonstrates how similar the two physical pictures are. By confining the waves in the cavity, light reflected from the mirrors in the resonator can, under the right conditions, constructively interfere to greatly enhance the field amplitude. As a result, an atom placed in the resonator light field couples much more strongly to it than to free space fields. Moreover, when these right conditions are not met, the light reflected from the mirrors destructively interferes, and the coupling is actually weaker than to free space.

The other effect of the resonator’s confinement is to modify the density of radiation modes that the atom can emit into the resonator compared to into free space. Following either a classical or quantum analysis, one finds that the resonator modifies the spontaneous emission rate of an excited atom: this is known as the Purcell effect [38, 39]. Under the constructive interference condition, the mode density in the resonator is larger than in free

space, stimulating an excited atom to emit light at a much larger rate into the resonator: this is Purcell enhancement, and we can use it to channel nearly all of the atom's emission into the resonator. Under the destructive interference condition, the mode density may be vanishingly small, suppressing the atom's emission into the resonator: this is Purcell suppression.

Altogether, we may utilise both the strong coupling of the atom to the resonator field and the Purcell effect to force single atoms to exchange their energy with single cavity photons that can travel in/out from an optical fibre, therefore constituting an atom-photon quantum interface between a stationary qubit (atom) and travelling qubit (photon).

With the above in mind, we define the framework for these atom-cavity interactions in this chapter: cavity quantum electrodynamics (CQED). Firstly in section 2.1, we define the Fabry-Pérot cavity (FPC) and its properties,¹ so that we can find the structure of the cavity mode field. Then, throughout section 2.2, we explore the 'quantum' in CQED. Starting in section 2.2.1, we quantise the field, using this to describe the coupling of a two-level atom to the cavity field.² Then in section 2.2.2, we show how the cavity channels the atom's emission into it using the Purcell effect, and how losses affect this. In section 2.2.3, we gather all of these effects and extend the model to a three-level atom to demonstrate how to use the system as a deterministic source of single photons. Finally, we summarise this chapter in section 2.3.

2.1 Optical cavity

An optical cavity, or resonator, can both 'pick out' a frequency of interest—such as an atomic transition frequency—and effectively multiply the intensity of light within by a large amount, thereby making the interaction between an atom and the light much stronger. We want to know

¹ Only FPCs will be used in this manuscript, so hereafter, 'cavity' and FPC will be used interchangeably.

² In the following, it will usually be sufficient to consider the field of a single mode as independent from any other mode. Only one mode of the field will be considered at a time, so hereafter, 'mode', 'field', and 'mode field' will be used interchangeably.

how much stronger the interaction becomes. So, in this section, the goal is to calculate the light field in the resonator, for use in calculating the atom-light coupling later in section 2.2.

2.1.1 Fabry-Pérot resonator

Suppose light of angular frequency ω and amplitude E_{in} is incident upon a Fabry-Pérot resonator of length L_{cav} , as depicted in figure 2.2a. Two lossless mirrors of reflection amplitudes r_1 and r_2 and corresponding transmission amplitudes t_1 and t_2 are aligned facing each other.³ Inside the cavity, the in-phase components of successive wavefronts constructively interfere. After each round-trip, E_{in} interferes with the circulating light that has already made one round-trip inside the cavity, E_{circ} , thus giving the recursion relation

$$E_{\text{circ}} = e^{i\delta} r_1 r_2 E_{\text{circ}} + t_1 E_{\text{in}}, \quad (2.1)$$

where

$$\delta = \frac{2\omega L_{\text{cav}}}{c_0}$$

for speed of light in vacuum c_0 is the phase picked up after one round-trip, as can be seen from figure 2.2a. So, the transmitted amplitude is

$$E_t = t_2 E_{\text{circ}}. \quad (2.2)$$

Then, substituting equation 2.1, the steady-state amplitude after many successive round-trips is

$$E_t = \frac{t_1 t_2}{1 - r_1 r_2 e^{i\delta}} E_{\text{in}}. \quad (2.3)$$

³ Note the distinction between amplitudes and intensities, the latter of which are $R = r^2$ and $T = t^2$. In particular for a lossless mirror, $r^2 + t^2 = R + T = 1$.

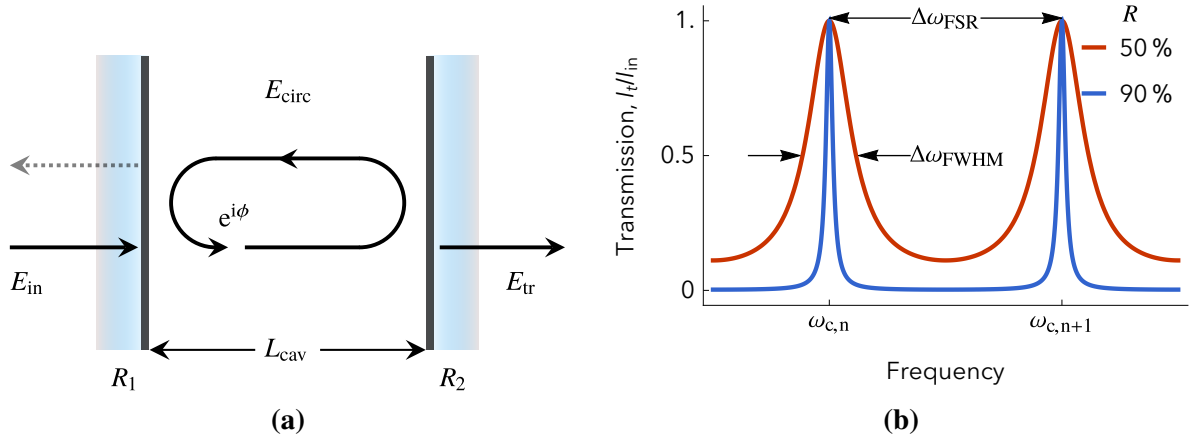


Figure 2.2: Properties of a driven FPC. **(a)** Light field transmission and reflection inside the cavity. After each cavity round-trip, the light field picks up an additional phase component ϕ . **(b)** Transmitted intensity as a function of the input frequency. A Lorentzian lineshape arises for the resonant cavity transmission after many round-trips (see equation 2.4), due to the constructive interference of the successive wavefronts.

Taking the modulus squared gives the transmitted intensity, $I_t = |E_t|^2$, which we rearrange⁴ to get

$$I_t = \left(\frac{t_1 t_2}{1 - r_1 r_2} \right)^2 \frac{I_{in}}{1 + \frac{4\mathcal{F}^2}{\pi^2} \sin^2\left(\frac{\delta}{2}\right)}, \quad (2.4)$$

with the finesse defined as

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2}. \quad (2.5)$$

So, the more reflective the mirrors, the higher \mathcal{F} is.

Note that equation 2.4 is maximised when

$$\frac{\delta}{2} = \frac{\omega L_{cav}}{c_0} = n\pi, \quad n \in 1, 2, 3 \dots$$

which is satisfied for the resonant frequencies of the cavity,

$$\omega_c = n\Delta\omega_{FSR} = \frac{n\pi c_0}{L_{cav}}. \quad (2.6)$$

⁴ Using $\cos \delta = (e^{i\delta} + e^{-i\delta})/2$ and $2 \sin^2(\delta/2) = 1 - \cos \delta$.

$\Delta\omega_{\text{FSR}}$ is known as the free spectral range. Thus, each ω_c corresponds to standing-wave resonances separated by half-wavelengths, or by $\Delta\omega_{\text{FSR}}$ in frequency, as illustrated in figure 2.2b.

We can link \mathcal{F} to the linewidth, ω_{FWHM} (**full-width at half-maximum**), corresponding to each ω_c using (see e.g. [40] for a derivation)⁵

$$\mathcal{F} = \frac{\Delta\omega_{\text{FSR}}}{\Delta\omega_{\text{FWHM}}}. \quad (2.7)$$

Accordingly, \mathcal{F} is a measure of how sharp the resonance is, and therefore with what resolution we may ‘pick out’ a particular frequency. Physically, \mathcal{F} also represents the \mathcal{F}/π enhancement of I_{circ} at resonance, as equation 2.4 shows for the ratio $I_{\text{circ}}/I_{\text{in}}$.⁶ In particular, at a standing-wave antinode, the in-phase addition of the field contributions from the left and right double the amplitude, and therefore we have $4\mathcal{F}/\pi$ enhancement at an antinode.⁷ Putting this all together, we now see the advantage that the cavity provides: a high \mathcal{F} greatly increases the intracavity light intensity whilst greatly restricting the possible frequencies that the light can take. In particular, these properties can be exploited to either enhance or inhibit the spontaneous emission of an atom into the cavity mode using the Purcell effect [38]. As a prerequisite for examining this effect later on, we will now calculate the structure of the field inside the resonator.

2.1.2 Resonator modes

What does the light field inside the cavity resonator look like? At a high level, we can start by noting that the field on the cavity mirrors must vanish. Also, the mirrors ought to be concave,

⁵ This is valid when \mathcal{F} is high. When it is ~ 10 or lower, we must modify this expression slightly [41].

⁶ Far off resonance, I_{circ} is suppressed by $\approx 4\mathcal{F}/\pi$. We have assumed the mirrors are identical and highly reflective.

⁷ Similarly, far off resonance, we have \mathcal{F}/π suppression at a standing-wave antinode.

since this makes the cavity more stable with respect to length perturbations. Therefore, the light field inside the resonator has curved wavefronts whose curvature matches that of the mirrors. From this intuition, a precise expression for the field can be found by solving Maxwell's equations with this boundary condition. This results in, after making the paraxial approximation,⁸ the expression for the fundamental Gaussian mode of the resonator (derived in e. g. [42, 43]),

$$E_{00}(z, \rho) = \frac{E_0}{kq} \exp\left(ik\left(z + \frac{\rho^2}{2q}\right)\right), \quad (2.8)$$

for which the resulting shape of the mode is shown in figure 2.3. Here, E_{00} is the field amplitude of the fundamental transverse mode (labelled by subscript 00),⁹ k is the wavevector,¹⁰ E_0 is the maximum field amplitude, z is position on the cavity axis, and $\rho = \sqrt{x^2 + y^2}$ is transverse to the cavity axis. q is a co-ordinate substitution,

$$q = z - iz_R, \quad (2.9)$$

whose meaning is soon made clear.

What is immediately apparent from equation 2.8 and figure 2.3 is the Gaussian intensity profile transverse to the cavity axis. In addition, because it is a standing wave, the mode has nodes and antinodes. The standing wave's antinodes have intensity $I_{\max} = (4\mathcal{F}/\pi) I_{\text{in}}$, following equation 2.4; the nodes have $I = 0$. It is therefore paramount that we must position atoms close to an antinode and close to the cavity axis for the strongest possible interaction with the field.

We desire an intuitive way to describe the strength of the field at a particular position. For

⁸ The paraxial approximation states that the wave does not diverge by large angles from its axis of propagation, and so the vectorial component of the field transverse to the cavity axis can be neglected.

⁹ We normalised E_{00} such that the beam carries a power $P = 2\pi \int |E_{00}(z, \rho)|^2 \rho d\rho$, which gives $2P/(c_0 \varepsilon_0) = \pi w_0^2 E_0^2$ [42]. ε_0 is the electric constant and w_0 is the beam waist, which we will soon define.

¹⁰ Wavevector is a valid name here; see [42] for a detailed justification of why we can ignore the direction of \mathbf{k} in our analysis.




Figure 2.3: Shape of the fundamental Gaussian mode in a cavity resonator. **(a)** Illustration of how the wave is spherical close to the mirrors, and then converges to a plane wave at the origin $z = 0$ where the beam waist is. In the Rayleigh zone (shaded), the mode is highly focused. The beam diameter doubles at z_R from the origin. **(b)** Photo showing the intensity of the same mode (reproduced from [42]). The nodes of zero intensity and antinodes of maximum intensity can be seen, as well as the Gaussian transverse profile and reduced intensity away from the beam waist.

this, we define the Rayleigh length as

$$z_R = \frac{\pi w_0^2}{\lambda} \quad (2.10)$$

for wavelength λ , and beam radius

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}. \quad (2.11)$$

The physical meanings of equations 2.9, 2.10 and 2.11 are therefore made clear:

1. q is a co-ordinate substitution that sets the wave to behave like a plane wave at $z = 0$ (hereafter referred to as the origin).
2. z_R is the length for which the beam diameter doubles from the origin. Past the Rayleigh zone $z = \pm z_R$, the wave behaves like a spherical wave that is emanating from the origin.
3. w_0 is the beam waist radius, where the beam takes its smallest radius and whose position is at the origin $z = 0$.

From this, we can see that there is an effective focal point of the resonator at the origin, and that close to it (i.e. within z_R), the beam divergence changes quickly. In comparison, at large z , the beam maintains a constant divergence. This scenario is just like that of a laser beam focused by a lens. So, the strongest field is found at this focus, but the strength quickly tails off and then exhibits a $1/\rho$ dependence further away. Therefore, an atom should ideally be placed at this focal point.

Completing the definitions, note that w_0 for an FPC is a function of the cavity geometry [44],

$$w_0^4 = \frac{\lambda^2 L_{\text{cav}} (\zeta_1 - L_{\text{cav}}) (\zeta_2 - L_{\text{cav}}) (-L_{\text{cav}} + \zeta_1 + \zeta_2)}{\pi^2 (\zeta_1 + \zeta_2 - 2L_{\text{cav}})^2}, \quad (2.12)$$

where ζ_1, ζ_2 are the radii of curvature of the two mirrors.¹¹ Since a smaller w_0 corresponds to a more strongly focused beam, ζ_1 and ζ_2 should be as small as possible to increase the intensity at the focus, giving a stronger interaction with an atom. Furthermore, we define the distance from the left mirror to the origin/beam waist as

$$z_0 = \frac{L_{\text{cav}}(\zeta_2 - L_{\text{cav}})}{\zeta_1 + \zeta_2 - 2L_{\text{cav}}}, \quad (2.13)$$

which shows that if $\zeta_1 = \zeta_2$, the beam waist is located precisely halfway between the mirrors, giving the best optical access to the cavity mode. Otherwise, the beam waist is closer to whichever mirror is flatter.

With all of the definitions now made, we make the structure of the field more clear by rewriting equation 2.8 as [42, 45]

$$E_{00}(z, \rho) = E_0 \frac{w_0}{w(z)} \exp \left(- \left(\frac{\rho}{w(z)} \right)^2 + i \left(\frac{k\rho^2}{2\zeta(z)} + kz - \alpha_G(z) \right) \right). \quad (2.14)$$

¹¹ For reference, if $\zeta_1 = \zeta_2 = \zeta$, then this equation simplifies to $w_0 = \sqrt{\frac{\lambda}{2\pi}} (L_{\text{cav}}(2\zeta - L_{\text{cav}}))^{1/4}$.

We can now summarise the salient features:

- $w_0/w(z)$ describes the mode's decreasing amplitude with increasing distance from the origin.
- $-\left(\frac{\rho}{w(z)}\right)^2$ in the exponential gives the Gaussian transverse distribution of the light. At distance $\rho = w$ from the cavity axis, the beam amplitude is reduced to $1/e \approx 37\%$ of its value on-axis.
- $\zeta(z) = z + \frac{z_R^2}{z}$ is the radius of the wavefront.
- e^{ikz} gives the phase evolution along z .
- The final term, α_G , is the Gouy phase [42].¹² It sets the phase to zero at the origin. Practically, its effect is to modify ω_c slightly, and as such will be considered shortly.

Although we have shown the full form of the resonator mode, there are two additional considerations when determining how strongly an atom couples to the cavity mode.

Mode volume

Firstly, the interaction strength is dependent on the cavity's mode volume, V_m . The smaller V_m is, the stronger the interaction, because the field strength increases as V_m decreases. This can be thought of as light bouncing back and forth within a smaller space, and therefore being more likely to pass through the absorption cross-section of the atom, $\sigma_{\text{abs}} = 3\lambda^2/(2\pi)$, where λ is the wavelength of the atomic transition.¹³ We can calculate V_m by integrating the field

¹² The Gouy phase results from the phase shift of π for a beam passing through a focus: we can see this in figure 2.3, where the curvature of the wavefront on one side of the mode waist flips direction on the other side.

¹³ When both the light focusing and the atom-light interaction are strong, one must modify σ_{abs} accordingly. We may interpret this as modifying how the incident and scattered light fields interfere, changing the area in which the atom can absorb strongly, due to how the atomic absorber changes the phase of the focused light after the interaction [46]. For us, this effect will manifest itself in the Purcell modification of the emission rate, so we will not need to consider it explicitly.

over the cavity volume. However, it can be more useful to define an effective mode volume, V_{eff} , that is dependent only on the cavity geometry and not the amplitude E_{00} . To do this, we divide V_m by the (maximal) field strength at the mode waist [47],

$$V_{\text{eff}} = \frac{V_m}{E_{00}(0,0)^2} = \frac{\int_{\text{cav}} |E_{00}(z, \rho)|^2 \rho \, d\rho \, d\phi \, dz}{E_{00}(0,0)^2}. \quad (2.15)$$

For reasons that will be described in chapter 4, the cavities we construct will target $\zeta_1 = \zeta_2$ and $L_{\text{cav}} = \frac{\zeta}{2}$. Inputting these conditions into equation 2.14, V_{eff} then takes the form

$$V_{\text{eff}} = \frac{\pi w_0^2 L_{\text{cav}}}{2}, \quad (2.16)$$

highlighting that smaller ζ and L_{cav} will increase the interaction strength.

Higher-order modes

Secondly, there are other possible transverse profiles that still solve Maxwell's equations and have a vanishing field on the mirror. These profiles can be described by Hermite polynomials; the corresponding modes are termed higher-order modes.¹⁴ The fundamental mode, E_{00} as in equation 2.14, is then multiplied by [45]

$$\mathcal{E}_{nm} = H_n\left(\sqrt{2}\frac{x}{w}\right) H_m\left(\sqrt{2}\frac{y}{w}\right), \quad (2.17)$$

where H_n and H_m are the Hermite polynomials of orders n and m corresponding to the modes along the x and y axes respectively. The higher the order of the mode, the larger α_G becomes, shifting the mode frequency accordingly. The frequency difference between mode orders

¹⁴ There are e. g. Laguerre-Gaussian modes that correspond to a cylindrical symmetry, but those do not apply to the fibre-tip mirrors described in this thesis, which exhibit an elliptical symmetry (see chapter 4) and therefore Hermite-Gaussian modes.

is [43]

$$\Delta\omega_{\text{trans}} = \frac{\Delta\omega_{\text{FSR}}}{\pi} \cos^{-1} \left(\sqrt{\frac{(\zeta_1 - L_{\text{cav}})(\zeta_2 - L_{\text{cav}})}{\zeta_1 \zeta_2}} \right). \quad (2.18)$$

In our experiments, higher-order modes for a well-aligned cavity tend to be lossier than the 00 mode,¹⁵ so we will usually only consider the fundamental mode. Nevertheless, the regular spacing between mode orders can be useful in characterising the cavity, as shall be seen in chapter 4.

To summarise, we have found the structure of the cavity field, and with it the knowledge of where to place the atom for maximum coupling. Therefore, we proceed by introducing an atom into the cavity.

2.2 Atom-cavity interactions

To set up a physical picture of atom-cavity interactions, let us first use some intuition to extend our atom-light interaction model to an atom-photon model. We described previously how the cavity confines the light to a small mode volume V_m , increasing its field strength and making it more likely to pass through the atom's absorption cross-section, and thus making atom-light interactions more likely. Similarly, a photon in the cavity will bounce back and forth between the mirrors, passing and therefore interacting with the atom \mathcal{F}/π times more than in free space. So, we would expect the atom-photon interaction in the cavity to be stronger. However, we must be cautious: we have to consider the difference in volume occupied by free space and cavity photons respectively.

Now, at the quantum level, we must draw a correspondence between the classical and quantum light modes if their results are to be consistent with each other. In doing so, one

¹⁵ This is because, for our fibre cavity, the light must hit the centre of the cavity mirror in order to reach the fibre core, and higher-order cavity modes are much less intense close to this centre.

finds that the amplitude of a single photon $A \propto 1/\sqrt{V_m}$ [48]. What should V_m be to maintain this consistency? Intuitively, free space can be thought of as a very, or even infinitely, large fictitious cavity, with similarly large V_m . To make this practical, we may consider in our quantisation procedure that the photon occupies only the region that we could possibly look for it in our experiment. For example, this may be the beam path from a laser on one side of a table to a detector on the other side. Then, V_m would be the size of this beam path. But, with our real optical cavity being much smaller, its V_m is smaller too, so the coupling strength of an atom to a single photon becomes much stronger than its coupling to a photon from this beam line and, by the same logic, its coupling to a free space photon. In this respect, the cavity can be said to enhance the atom-photon coupling well beyond the atom's coupling to a free space photon, both by allowing \mathcal{F}/π more opportunities for the atom and photon to interact, and by reducing V_m .

What we will now do is use **CQED** to precisely describe and quantify our intuition. In this section, we explain how the coupling strength of an atom to a cavity photon can become much stronger than its coupling to a free space photon, thereby reaching the goal of deterministic atom-photon coupling, as needed for a quantum interface.

2.2.1 Atom-cavity coupling

Two-level atom outside the cavity

Consider first an empty cavity, save for some light inside. The Hamiltonian of the its light field is

$$\hat{H}_c = \hbar\omega_c \left(\hat{c}^\dagger \hat{c} + \frac{1}{2} \right) = \hbar\omega_c \left(\hat{n} + \frac{1}{2} \right), \quad (2.19)$$

where \hbar is the reduced Planck constant, \hat{c}^\dagger and \hat{c} are the creation and annihilation operators for a cavity photon and $\hat{n} = \sum_{n=0}^{\infty} n|n\rangle\langle n|$ is the number operator, which has eigenvalue n if




Figure 2.4: Energy levels of the atom and cavity in the case $\omega_a = \omega_c$.

the field contains n photons.¹⁶

Separately, consider a two-level atom in free space, outside the cavity. In the basis $\{|g\rangle, |e\rangle\}$ the Hamiltonian is

$$\hat{H}_a = -\frac{\hbar\omega_a\hat{\sigma}_Z}{2}, \quad (2.20)$$

where $|g\rangle$ and $|e\rangle$ denote the ground and excited state levels respectively, $\omega_a = \omega_e - \omega_g$, and $\hat{\sigma}_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ is the Pauli-Z operator.¹⁷ Figure 2.4 depicts the energy levels of both the atom and the cavity.

Atom inside the cavity

When we place an atom in the cavity field, an additional term to describe the atom-cavity interaction, \hat{H}_{int} , arises. So far, the total Hamiltonian of the system is

$$\hat{H} = \hat{H}_a + \hat{H}_c + \hat{H}_{\text{int}}. \quad (2.21)$$

¹⁶ Note the inclusion of the zero-point energy, $\hbar\omega_c/2$.

¹⁷ Note that the energy of the atom is defined as zero halfway between $|g\rangle$ and $|e\rangle$.

\hat{H}_{int} describes an energy level shift due to the induced atomic dipole moment when the atom interacts with the cavity field. The classical energy of the induced dipole potential is¹⁸

$$H_{\text{int}} = -\mathbf{D} \cdot \mathbf{E} = e\mathbf{r} \cdot \mathbf{E}, \quad (2.22)$$

where \mathbf{D} is the atomic dipole moment, e is the electron charge and $\mathbf{E}(\mathbf{r})$ is the electric field at the atom's position,¹⁹ with \mathbf{r} being a position vector. What follows now is to quantise this expression, thereby arriving at a simple form that fully encapsulates how the atom and cavity field interact at a quantum level.

To quantise equation 2.22, firstly we define the position operator of the valence electron involved,

$$\begin{aligned} \hat{r} &= \hat{\mathbb{I}} \hat{r} \hat{\mathbb{I}}, \\ &= \langle g|\hat{r}|g\rangle |g\rangle\langle g| + \langle g|\hat{r}|e\rangle |g\rangle\langle e| + \langle e|\hat{r}|g\rangle |e\rangle\langle g| + \langle e|\hat{r}|e\rangle |e\rangle\langle e|, \\ &= \langle g|\hat{r}|e\rangle \hat{\sigma}^- + \langle e|\hat{r}|g\rangle \hat{\sigma}^+, \\ &= \mathbf{r}_{\text{ge}} \hat{\sigma}^- + \mathbf{r}_{\text{ge}}^* \hat{\sigma}^+, \end{aligned} \quad (2.23)$$

where $\hat{\mathbb{I}}$ is the identity operator (which here is $|g\rangle\langle g| + |e\rangle\langle e|$), $\hat{\sigma}^+ = |e\rangle\langle g|$ and $\hat{\sigma}^- = |g\rangle\langle e|$ are raising and lowering operators for the atomic state, \mathbf{r}_{ge} denotes the electron's change in position for the particular transition, and the dipole approximation has been applied.²⁰ Note that \hat{r} is an odd-parity operator, so the first and fourth terms on the second line vanish.

Secondly, we must quantise the electric field, $\mathbf{E}(\mathbf{r})$. Proceeding with care, the first step is

¹⁸ An intuitive classical description of the dipole interaction is available in section 2.1.2 of [49].

¹⁹ The interaction of the atom with the corresponding magnetic field is much weaker and can be ignored, see e. g. [48].

²⁰ This is where $\lambda = 2\pi c/\omega_a$ is much larger than the atom, such that $\mathbf{E}(\mathbf{r})$ is roughly constant over it.

to write the fundamental Gaussian mode as defined in equation 2.14 as the product of the ‘base’ amplitude of the field and the position-dependent spatial part,

$$\mathbf{E}_{00}(z, \rho) = \mathbf{E}_0 E_r(z, \rho). \quad (2.24)$$

Immediately, we recognise that E_r is maximal at $(z = 0, \rho = 0)$, i. e. antinode at the mode waist. Here, $E_r(0, 0) = 1$. For the following, we shall assume that the atom is at this position of maximal coupling.

With the position dependence taken care of, next we must find the value of \mathbf{E}_0 . Its value naturally arises during the quantisation process of the EM field. The standard textbook approach to this quantisation is to solve Maxwell’s equations in the cavity, again using the boundary condition that the field on the mirrors must vanish [48, 50]. This confines the solution to the volume of the cavity mode, V_m . Assuming that the cavity only supports a single mode at a time, and considering only one particular polarisation of the field, then the energy inside the EM field is

$$H_{\text{EM}} = \frac{1}{2} \int_V \left(\varepsilon_0 |\mathbf{E}(\mathbf{r})|^2 + \frac{|\mathbf{B}(\mathbf{r})|^2}{\mu_0} \right) dV = 2\varepsilon_0 V_m |\mathbf{E}_0|^2, \quad (2.25)$$

where ε_0 and μ_0 are the electric and magnetic constants respectively.²¹ Let us relate this to the energy of a single cavity photon, which has energy $\hbar\omega_c$:

$$2\varepsilon_0 V_m |\mathbf{E}_0|^2 = \hbar\omega_c, \quad (2.26)$$

$$\therefore \mathbf{E}_0 = \hat{\mathbf{e}} \sqrt{\frac{\hbar\omega_c}{2\varepsilon_0 V_m}}, \quad (2.27)$$

where the direction of \mathbf{E}_0 (and \mathbf{E}) is simply the direction of its polarisation, denoted explicitly

²¹ Also referred to as vacuum permittivity and permeability respectively.

here as the unit vector $\hat{\mathbf{e}}$.

Lastly, we quantise the classical electric field by expressing its observables in terms of the cavity photon annihilation and creation operators \hat{c} and \hat{c}^\dagger . By comparing equation 2.26 with equation 2.19, we note that the classical energy corresponds to one step of a quantum harmonic oscillator. Thus, we can relate the energy eigenstates of the quantum harmonic oscillator with the classical energy. At the end of the textbook quantisation process, this relation is used to define the electric field observable, which at the position of maximum coupling is

$$\hat{\mathbf{E}} = \mathbf{E}_0 (\hat{c} + \hat{c}^\dagger). \quad (2.28)$$

We can now quantise H_{int} . Substituting equations 2.23 and 2.28 into equation 2.22 gives

$$\begin{aligned} \hat{H}_{\text{int}} &= e\mathbf{E}_0 \cdot (\mathbf{r}_{\text{ge}}\hat{\sigma}^- + \mathbf{r}_{\text{ge}}^*\hat{\sigma}^+) (\hat{c} + \hat{c}^\dagger), \\ &= eE_0 r_{\text{ge}} (\hat{\sigma}^- + \hat{\sigma}^+) (\hat{c} + \hat{c}^\dagger), \end{aligned} \quad (2.29)$$

where we have dropped the vector notation, and used $r_{\text{ge}} = r_{\text{ge}}^*$ by assuming the polarisations of the light and the atomic transition are the same.

We want to focus on this interaction now. To do so, we look for a simple form of the total Hamiltonian, equation 2.21, that unclutters the uncoupled atom-cavity system dynamics from the coupled system. From the Schrödinger equation,

$$\frac{d|\psi(t)\rangle}{dt} = -\frac{i}{\hbar}\hat{H}|\psi(t)\rangle, \quad (2.30)$$

where $|\psi(t)\rangle$ is the system state at time t , we find the evolution operator,

$$\hat{U}(t) = e^{-\frac{i}{\hbar}\hat{H}t}, \quad (2.31)$$

that evolves the system state from $|\psi(0)\rangle$ to $|\psi(t)\rangle$. Let us use this to transform \hat{H} into the interaction picture by setting an interaction Hamiltonian [51]

$$\hat{H}_I = \hat{U}_0^\dagger \left(\hat{H} - \hat{H}_a - \hat{H}_c \right) \hat{U}_0 = \hat{U}_0^\dagger \hat{H}_{\text{int}} \hat{U}_0, \quad (2.32)$$

where \hat{U}_0 is the evolution operator corresponding to just \hat{H}_a and \hat{H}_c ,²²

$$\hat{U}_0 = \hat{U}_a \hat{U}_c.$$

We have assumed that the cavity is on resonance with the atom, $\omega_c = \omega_a = \omega$. Then, using the spectral theorem,²³ we have for a single cavity photon

$$\begin{aligned} \hat{U}_a &= e^{i\frac{\omega}{2}t} |g\rangle\langle g| + e^{-i\frac{\omega}{2}t} |e\rangle\langle e|, \\ \hat{U}_c &= |0\rangle\langle 0| + e^{-i\omega t} |1\rangle\langle 1|. \end{aligned}$$

After a few more lines and applying the rotating wave approximation (RWA),²⁴ we get

$$\begin{aligned} \hat{H}_I &= eE_0 r_{\text{ge}} \left(\hat{\sigma}^- e^{-i\omega t} + \hat{\sigma}^+ e^{i\omega t} \right) \left(\hat{c} e^{-i\omega t} + \hat{c}^\dagger e^{i\omega t} \right) \\ &\stackrel{\text{RWA}}{\implies} = \hbar g (\hat{\sigma}^+ \hat{c} + \hat{\sigma}^- \hat{c}^\dagger), \end{aligned} \quad (2.33)$$

where g is the atom-cavity coupling parameter,

$$g = \frac{DE_0}{\hbar} = e r_{\text{ge}} \sqrt{\frac{\omega_c}{2\hbar\epsilon_0 V_m}}. \quad (2.34)$$

²² For clarity, we have omitted both the time dependence and the tensor product notation ($\hat{U}_a(t) \otimes \hat{U}_c(t)$).

²³ In quantum physics, we can state this as $f(\hat{A}) = \sum_{n=1}^N f(\lambda_n) |\lambda_n\rangle\langle\lambda_n|$, where $f(\hat{A})$ is a function of some observable \hat{A} , and λ_n and $|\lambda_n\rangle$ are its eigenvalues and eigenvectors.

²⁴ This states that if $|\omega_a - \omega_c| \ll \omega_a + \omega_c$, then the fast-oscillating exponential terms with $\omega_a + \omega_c$ can be neglected, since the timescale of the interaction $\gg 1/(\omega_a + \omega_c)$ and so has negligible effect on average.

These equations give the key insight that the interaction causes the atom and cavity to continuously exchange a single quantum of energy. The first term in equation 2.33 creates an atomic excitation while destroying a cavity photon; the second term creates a cavity photon whilst de-exciting the atom. A larger g increases the rate of this exchange. If there were no losses, this exchange would carry on forever at a Rabi frequency of $2g$, and therefore total control of g would allow total control of the atom-cavity interaction and this exchange (cf. Rabi laser pulses for ions). Later in section 2.2.2, we will consider how to design the cavities to increase g in spite of the losses in realistic experiments.

Vacuum Rabi oscillations

What precisely is the nature of the exchange? Being at a quantum level, we might expect it to happen in a discrete manner, with the excitation instantaneously swapping between atom and cavity. Our intuition would tell us this is not the case—equation 2.33 resembles a continuous cycle of stimulated emission and absorption. A precise answer follows from finding the evolution of the system.

We proceed by calculating the energy eigenstates of \hat{H}_I (equation 2.33) for a single excitation,

$$\begin{aligned} |+, 1\rangle &= \frac{1}{\sqrt{2}} (|g, 1\rangle + |e, 0\rangle), \\ |-, 1\rangle &= \frac{1}{\sqrt{2}} (|g, 1\rangle - |e, 0\rangle), \end{aligned} \tag{2.35}$$

where $|g, n\rangle = |g\rangle \otimes |n\rangle$: that is, the atom-cavity state is written as a tensor product of atomic state and cavity state. These eigenstates have energy eigenfrequencies $\pm g$. Therefore, these states are separated in frequency by $2g$.

Now, suppose the state of the atom-cavity system starts at time $t = 0$ in $|\psi(0)\rangle = |g, 1\rangle$. Still remaining in the interaction picture, then by applying Schrödinger's equation 2.30 and

the spectral theorem to \hat{H}_I , and using equation 2.35, we get

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}} \left(e^{igt} |+, 1\rangle + e^{-igt} |-, 1\rangle \right), \quad (2.36)$$

which can be transformed back to the previous basis by again using equation 2.35, giving

$$|\psi(t)\rangle = \cos(gt) |g, 1\rangle + \sin(gt) |e, 0\rangle. \quad (2.37)$$

Immediately, we see that the system oscillates between $|g, 1\rangle$ and $|e, 0\rangle$, according to

$$\begin{aligned} P_{|g,1\rangle}(t) &= \cos^2(gt), \\ P_{|e,0\rangle}(t) &= \sin^2(gt), \end{aligned} \quad (2.38)$$

where $P(t)$ is the probability of being in a particular state at time t . Remarkably, the interaction is just like that of Rabi oscillations of a laser-driven atom in free space—except here, there is no laser. Even in a cavity with zero photons, the fluctuations of the cavity's vacuum field will stimulate an excited atom in the cavity to oscillate between $|g, 1\rangle$ and $|e, 0\rangle$ at Rabi frequency $2g$.²⁵ This phenomenon is therefore known as vacuum Rabi oscillation. Note that the cavity has picked out only a single mode of the vacuum. Because the atom only interacts with this single mode then, unlike spontaneous emission, the evolution of the system remains coherent,²⁶ and so crucially, we do not lose the quantum information contained in the system.

For this scenario, the associated splitting of the energies by $2\hbar g$ is known as vacuum Rabi splitting. Practically, this serves to split a single cavity resonance peak into two resonant

²⁵ To clarify the factor of 2 difference between frequency and Rabi frequency: for these equations, the initial state matches the evolving state whenever $t = \pi/g$. However, the Rabi frequency Ω is defined as the frequency for which an evolution of length $t = 2\pi/\Omega$ will match the initial state. Therefore, the Rabi frequency is $2g$.

²⁶ Spontaneous emission itself is stimulated emission caused by vacuum fluctuations in free space. Practically, this is an incoherent process because it is impossible to measure the phase/coherence between the state of the (atomic) system and the state of the (vacuum) environment that contains all possible modes.

peaks. The larger g is, the greater the splitting, and we can directly observe this in the laboratory by scanning the cavity's resonant frequency and the atomic transition frequency near these resonances. Although, it is not quite that simple, as shall be seen later when considering $\omega_c \neq \omega_a$.

N photons, resonant case

What happens if there is more than one cavity photon? We can answer this by applying the 'ladder' property of the cavity photon operators,

$$\begin{aligned}\hat{c}|n\rangle &= \sqrt{n}|n-1\rangle, \\ \hat{c}^\dagger|n\rangle &= \sqrt{n+1}|n+1\rangle,\end{aligned}\tag{2.39}$$

to \hat{H}_I (equation 2.33), giving the following relations for n quanta in the system:

$$\begin{aligned}\hat{H}_I^{(n)}|g, n\rangle &= \hbar g \sqrt{n}|e, n-1\rangle, \\ \hat{H}_I^{(n)}|e, n-1\rangle &= \hbar g \sqrt{n}|g, n\rangle.\end{aligned}\tag{2.40}$$

Here, the (n) superscript denotes the Hamiltonian that is restricted only to the case of n quanta. From these relations, we immediately see that the interaction energy scales as \sqrt{n} . Therefore, we can repeat the earlier analysis simply by replacing $g \rightarrow g\sqrt{n}$.

What is the total energy of the system? It is equivalent to that of the uncoupled system (atom outside cavity), $\hat{H}_a + \hat{H}_c$, plus that of the atom-cavity interaction, \hat{H}_{int} . Luckily, we have already found the interaction energy as $\pm\hbar g$ (or $\pm\hbar g\sqrt{n}$), since this precisely corresponds to the eigenenergies of \hat{H}_I (equation 2.35): this is because \hat{H}_I is just \hat{H}_{int} after diagonalisation. Hence, for n quanta in the system and m cavity photons, the total energy of the coupled

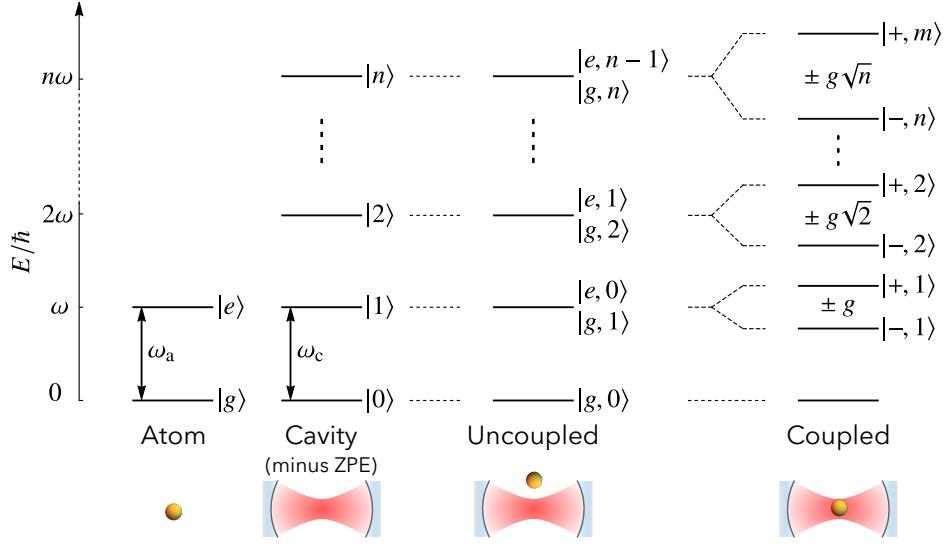


Figure 2.5: The Jaynes-Cummings ladder when $\omega_a = \omega_c$. When the atom is coupled to the cavity, the energy levels of the atom-cavity system are split into a series of doublets. For clarity, we have omitted the energy offsets of the bare atom by $-\hbar\omega_a/2$ and bare cavity by the zero-point energy (ZPE) $\hbar\omega_c/2$.

system in an energy eigenstate is

$$\begin{aligned} \frac{E_{\text{tot}}}{\hbar} &= \pm \frac{\omega_a}{2} + \omega_c \left(m + \frac{1}{2} \right) \pm g\sqrt{n} \\ &= n\omega \pm g\sqrt{n} \end{aligned} \quad (2.41)$$

in the resonant case of $\omega_a = \omega_c$, and where we have also set the value of m to conserve the energy of the system: if the atom is in state $|g\rangle$, meaning that the energy of the atom is $E_a/\hbar = -\omega/2$, then $m = n$; or if in $|e\rangle$ and $E_a/\hbar = +\omega/2$, then $m = n - 1$.

Overall, the effect of the atom-cavity interaction is to ‘dress’ the energy eigenstates of the uncoupled system, giving rise to a ladder of doublets (apart from the $|g, 0\rangle$ level, which is not split) known as the Jaynes-Cummings ladder, illustrated in figure 2.5.

Off-resonant case

What happens when $\omega_c \neq \omega_a$? Then, the interaction picture used in equation 2.32 is different: the atom is being driven off-resonantly by the cavity at $\omega_c \neq \omega_a$, so instead of the $\hat{H} - \hat{H}_a - \hat{H}_c$ term, we instead have²⁷

$$\hat{H}_{I,\Delta} = \hat{U}_0^\dagger \left(\hat{H} - \hat{H}_{a,\Delta} - \hat{H}_c \right) \hat{U}_0 = \hat{U}_0^\dagger \left(-\frac{\hbar\Delta\hat{\sigma}_Z}{2} + \hat{H}_{\text{int}} \right) \hat{U}_0, \quad (2.42)$$

where

$$\begin{aligned} \Delta &= \omega_a - \omega_c, \\ \hat{H}_{a,\Delta} &= -\frac{\hbar\omega_c\hat{\sigma}_Z}{2}. \end{aligned}$$

Repeating the calculation with $\omega_c = \omega$ yields

$$\stackrel{\text{RWA}}{\implies} \hat{H}_{I,\Delta} = -\frac{\hbar\Delta\hat{\sigma}_Z}{2} + \hbar g(\hat{\sigma}^+ \hat{c} + \hat{\sigma}^- \hat{c}^\dagger). \quad (2.43)$$

The new energy eigenfrequencies and eigenstates for this off-resonant interaction are, after some calculation,

$$\omega_\Delta = \pm \frac{1}{2} \sqrt{\Delta^2 + 4g^2}, \quad (2.44)$$

$$\begin{aligned} |+, 1\rangle &= \left(\frac{-\Delta - \sqrt{\Delta^2 + 4g^2}}{2g} \right) |g, 1\rangle + |e, 0\rangle, \\ |-, 1\rangle &= \left(\frac{-\Delta + \sqrt{\Delta^2 + 4g^2}}{2g} \right) |g, 1\rangle + |e, 0\rangle. \end{aligned} \quad (2.45)$$

²⁷ We consider the following for a single excitation $n = 1$, but as stated before, for larger n we can replace $g \rightarrow g\sqrt{n}$.

We can write the eigenstates in a more intuitive form by defining a mixing angle, θ , where (see appendix A)

$$|+, 1\rangle = \cos \theta |g, 1\rangle + \sin \theta |e, 0\rangle, \quad (2.46)$$

$$|-, 1\rangle = -\sin \theta |g, 1\rangle + \cos \theta |e, 0\rangle,$$

$$\tan \theta = \frac{-2g}{\Delta + \sqrt{\Delta^2 + 4g^2}}. \quad (2.47)$$

We also have the total energy for n quanta (cf. equation 2.41),

$$\begin{aligned} \frac{E_{\text{tot}}}{\hbar} &= \pm \frac{\omega_a}{2} + \omega_c \left(m + \frac{1}{2} \right) \pm \frac{1}{2} \sqrt{\Delta^2 + 4g^2 n}, \\ &= n\omega + \frac{\Delta \pm \sqrt{\Delta^2 + 4g^2 n}}{2}. \end{aligned} \quad (2.48)$$

The overall interaction therefore results in an avoided crossing, as displayed in figure 2.6. It is clear from the figure that if one tries to observe the split resonances of the atom-cavity system by scanning just the cavity, then the atom and cavity will uncouple, so this approach would be unsuccessful. Practically, another approach is to probe the system by driving the atom with a laser aimed through the open side of the cavity (i. e. oriented perpendicular to the cavity axis) and scan the laser frequency.²⁸ We will investigate this approach in chapter 5.

2.2.2 Losses: damped atom-cavity system

Up until now, losses from the atom-cavity system have not been under consideration. However, realistic cavities are lossy systems, open to interaction with their surrounding environment.

²⁸ It is possible to keep ω_c fixed and scan ω_a . Then, the same spectrum is produced, with the roles of atom and cavity reversed. However, this is experimentally more difficult for the systems considered in this manuscript. For example, the power of a laser could be modulated in order to Stark-shift ω_a [52]. Relating to our experimental schemes, we could use optical tweezers to achieve this [53].

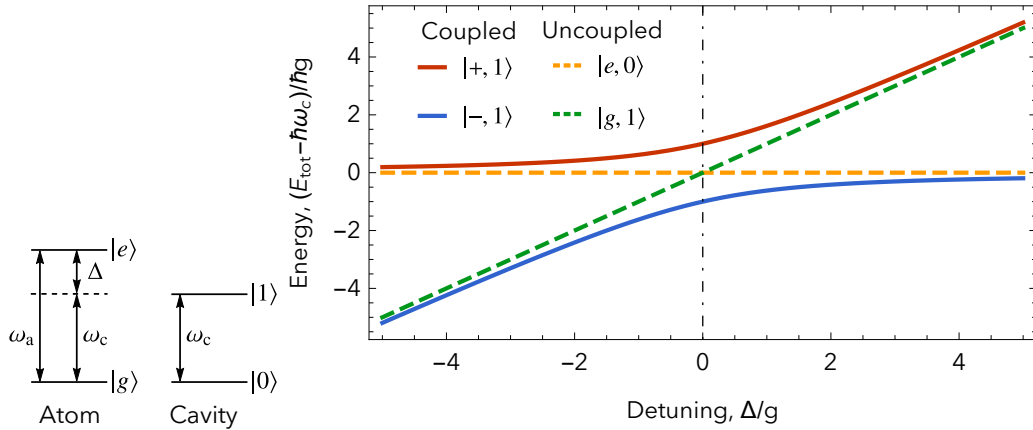


Figure 2.6: **Left:** Level structure of the two-level atom when the cavity is not in resonance with it. **Right:** Spectrum of the Jaynes-Cummings ladder's first rung when we scan the cavity frequency ω_c over the (fixed) atomic transition frequency ω_a . When uncoupled (dashed lines), changing ω_c —and therefore the energy corresponding to $|g, 1\rangle$ —has no effect on that of $|e, 0\rangle$. When coupled (solid lines), an avoided crossing results; the smaller the magnitude of Δ , the more the states $|\pm, 1\rangle$ are repelled by, up to a maximum of $2g$ on resonance.




Figure 2.7: Illustration of loss processes in an atom-cavity system. g , κ and γ are atom-cavity coupling strength, cavity field decay rate through the mirrors, and the atom's spontaneous (field) emission rate to the environment respectively.

So, the exchange of energy quanta between atom and cavity is hampered, with the Rabi oscillations accordingly damped by the losses, as we shall see. Figure 2.7 introduces the loss processes through the parameters κ and γ .

κ is the decay rate of the cavity field through the mirrors. Noting that the light intensity is $I \propto E^2$, and that the intensity decay rate of a non-driven high-finesse cavity is [40]

$$I = I_0(t)e^{-\Delta\omega_{\text{FWHM}}t},$$

then

$$2\kappa = \Delta\omega_{\text{FWHM}} = \frac{\Delta\omega_{\text{FSR}}}{\mathcal{F}}, \quad (2.49)$$

elucidating that for a fixed L_{cav} , a higher \mathcal{F} confines a cavity photon for longer before its inevitable escape. The decay rate of cavity photons, the number of which scales as the intensity, is therefore 2κ .

γ is the atomic amplitude decay rate. Correspondingly, 2γ is the atomic population decay rate of non-cavity photons to the environment.²⁹ More precisely, it is the rate of spontaneous emission into free space, where the atom goes from $|e\rangle$ to $|g\rangle$ due to interaction with the vacuum field. Here we see why γ counts as a loss: spontaneous emission is induced by fluctuations of the vacuum field, which is a continuum of all possible modes. Through this interaction, the environment is now a part of the atom's state, as well as the systems with which it is interacting coherently. Being unable to gain information about the environment's state, and therefore the system's interactions with the continuum of vacuum modes, we lose the coherences between the different parts of our system, and so spontaneous emission causes the atom to decohere.³⁰ γ can be defined by Fermi's golden rule [40, 50],

$$2\gamma = \frac{2\pi}{\hbar} \left| \mathbf{D}_{\text{eg}} \cdot \mathbf{E} \right|^2 \varrho(\omega), \quad (2.50)$$

where $\varrho(\omega)$ is the density of photon states, i.e. the available modes that the emitted photon can occupy over a particular frequency interval, and $\mathbf{D}_{\text{eg}} = -e\mathbf{r}_{\text{eg}}$. We can use this to see how the cavity is effectively able to channel the emission of the atom almost exclusively into the cavity mode.

²⁹ Thus, 2γ is also the atomic linewidth (full-width at half-maximum (FWHM)), using a similar approach as for equation 2.49.

³⁰ There are many ways to interpret this decoherence. Footnote 26 describes one way, wherein the impossibility of measuring the phases between the vacuum field fluctuations and the atom arises because we can only measure energy differences of a field from the ZPE and not the ZPE itself [50].

Purcell effect

Suppose an excited atom is in the cavity, like in figure 2.7. Then, it couples to the vacuum field. According to equation 2.50, the ratio of the emission into the cavity mode versus into a free space mode is

$$\frac{\gamma_c}{\gamma} = \frac{|\mathbf{D}_{eg} \cdot \mathbf{E}|_{cav}^2 \rho_{cav}(\omega)}{|\mathbf{D}_{eg} \cdot \mathbf{E}|_{free}^2 \rho_{free}(\omega)}.$$

Since emission is only along one direction in the cavity (effectively 1D), whereas in free space there are three distinct directions (3D), this results in $|\mathbf{D}_{eg} \cdot \mathbf{E}|_{cav}^2 / |\mathbf{D}_{eg} \cdot \mathbf{E}|_{free}^2 = 3$.³¹ Altogether, this gives the Purcell enhancement factor

$$F_P = \frac{3\rho_{cav}(\omega)}{\rho_{free}(\omega)}. \quad (2.51)$$

This makes calculating the enhancement factor quite simple.

The density of states in free space, over the volume of the mode V_m , is (without derivation, from [40])

$$\rho_{free}(\omega) = \frac{\omega^2 V_m}{\pi^2 c^3}. \quad (2.52)$$

In the cavity, we find the density of states by noting that the only mode available to emit into is the cavity mode, whose frequency has a Lorentzian lineshape.³² Thus, this gives (again over V_m and without derivation, from [40])³³

$$\rho_{cav}(\omega) = \frac{2}{\pi} \frac{\Delta\omega_{FWHM}}{4(\omega - \omega_c)^2 + \Delta\omega_{FWHM}^2}. \quad (2.53)$$

³¹ To see this, consider a unit sphere: a single unpolarised emission points $\hat{\mathbf{e}}$ to a random location on the sphere's surface. With respect to the cavity axis, let us set the projection angle θ and emission solid angle Ω . Thus, we have $\langle \cos^2 \theta \rangle_{\text{emiss}} = (1/4\pi) \oint_{\Omega} \cos^2 \theta \, d\Omega = (1/4\pi) \int_0^\pi \cos^2 \theta \sin \theta \, d\theta \int_0^{2\pi} d\phi = 1/3$. The reciprocal of this is our result.

³² See figure 2.2b and equation 2.4.

³³ If deriving this expression, the normalisation condition $\int_0^\infty \rho(\omega) \, d\omega = 1$ is used, together with a Fourier transform of the probability for a spontaneous emission as a function of time, where the emission is modelled as similar to a damped harmonic oscillator.

Altogether, on resonance and combining equations 2.51–2.53, we find that the Purcell factor is

$$F_P = \frac{3\mathcal{F}\lambda^2 L_{\text{cav}}}{2\pi^2 V_m}, \quad (2.54)$$

which is the ratio of emission into the cavity mode versus into a non-cavity mode. Of course, L_{cav} on resonance is an integer number of half-wavelengths, so F_P exhibits a λ^3 dependence. Remarkably, in a small cavity with highly reflective mirrors, this means almost all the emission is channelled into the cavity mode rather than lost to the environment/free space: very usefully, this inhibits loss of the photon to the environment during an atom-photon interfacing process. For our cavities, the fraction emitted into the cavity is close to unity.

Strong coupling

Now, we draw the connection between the atom-cavity coupling described previously in section 2.2.1 and the Purcell enhancement by defining the co-operativity [54],

$$C = \frac{F_P}{2} = \frac{g^2}{2\kappa\gamma}. \quad (2.55)$$

This new form clearly shows that C (and F_P) is the ratio of the coherent, desired atom-photon coupling to the incoherent, undesired losses from the system. From here, it is straightforward to see what Purcell enhancement is needed for a decent atom-photon interface.

Suppose $C < 1$. Then, the emission process is irreversible. Figure 2.8a shows how the system behaves like an overdamped oscillator, never quite able to return to its original state. The photon escapes through the mirrors or is lost to the environment before it has a chance to be re-absorbed by the atom. If $g \ll \kappa, \gamma$, this is the ‘weak coupling’ regime—not very useful. However, if $\kappa > g \gg \gamma$, then this is the ‘bad cavity’ regime, in which the atom-cavity coupling is strong but where the cavity photons rapidly leak out of the mirrors before a

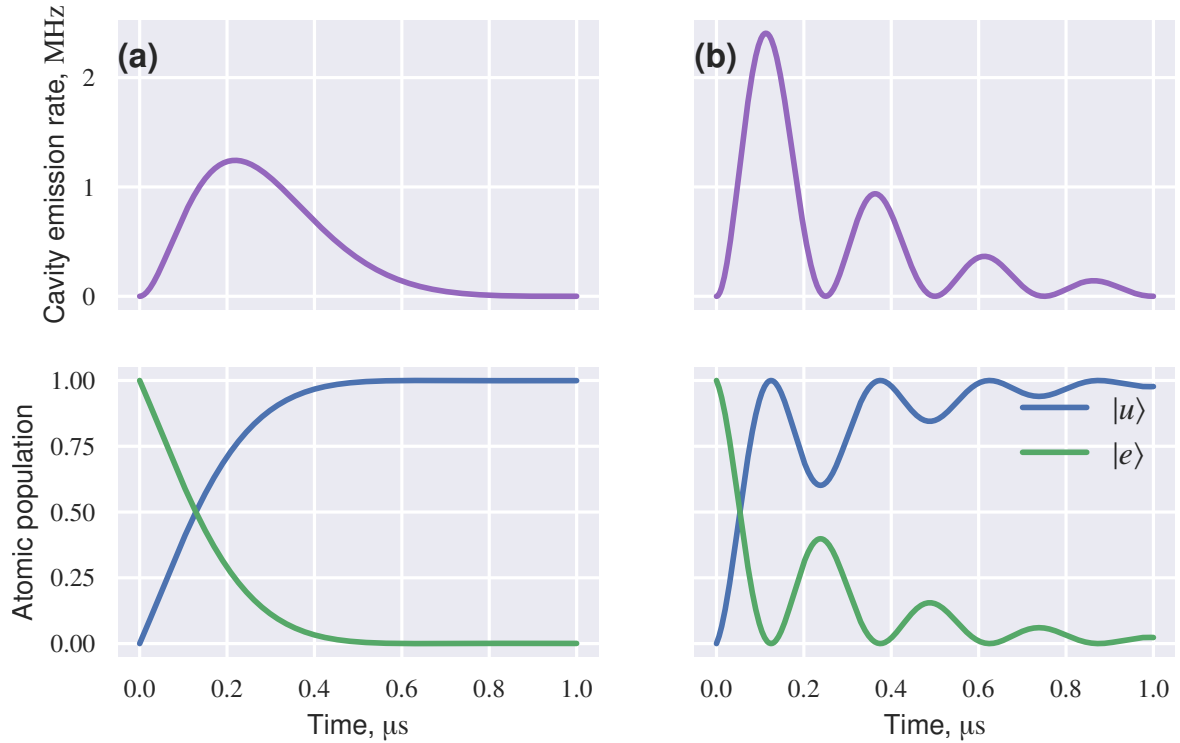


Figure 2.8: Atomic and cavity populations starting from $|e, 0\rangle$, illustrating how losses affect performance of an atom-cavity interface. **(a)** Bad cavity regime, with $(g, \kappa, \gamma)/2\pi = (0.6, 0.9, 0.3)$ MHz. **(b)** Strong coupling regime, with $(g, \kappa, \gamma)/2\pi = (2, 0.3, 0.3)$ MHz. Note that we can retrieve the cavity population from the emission rate simply by dividing the latter by 2κ . To plot these figures, we numerically solve a master equation, as we will describe in chapter 3.

coherent interaction can occur. The ‘bad cavity’ regime is still useful for some applications, since Purcell-enhanced emission still occurs.

Now, suppose $C > 1$. Then, a reversible exchange of a quantum between the cavity and atom becomes possible:

1. A single cavity photon excites the atom.
2. The atom emits a single photon into the cavity upon relaxation.
3. The emitted photon excites the atom once more.

Figure 2.8b shows that we retrieve the original atom-cavity state after one oscillation, albeit with a reduced population. This is the regime of strong coupling, allowing for the quantum

interface described in chapter 1. Because the evolution corresponding to the atom-cavity coupling is coherent and unitary, it is therefore also reversible. Thus, control of the atom-cavity interaction when in the strong coupling regime affords coherent and reversible control of the atom-cavity state. Altogether, with the high \mathcal{F} and small ζ of the fibre-tip Fabry-Pérot cavities used in this thesis, we can reach well into the strong coupling regime.

2.2.3 Single photon source

Three-level atom-cavity system

Up until now, we have used a two-level atom in our model. This is because, although a realistic atom has many levels, only the important levels need to be chosen: the two levels with transition frequency ω_a close to ω_c (ignoring for now levels with the same energies which may be coupled via different polarisations of light). But, if we choose three atomic levels, then Raman transitions³⁴ can be exploited to utilise the atom-cavity system as a deterministic, highly efficient single photon source. Similarly, the same process enables single photon absorption.

To demonstrate the single photon generation process, we extend the two-level atom-cavity system described in section 2.2.1 to a three-level system by defining a second ground state of the atom, $|u\rangle$. Figure 2.9 illustrates this Λ -level³⁵ system. With a suitable laser pulse, this Λ -system emits (absorbs) a single photon out into (out of) the cavity mode, thus forming the node in a quantum network of stationary and flying qubits just as described in figure 1.1.

Suppose we apply a laser pulse with frequency ω_L and time-dependent Rabi frequency $\Omega(t)$ to the atom on the $|u\rangle \rightarrow |e\rangle$ transition, as displayed in figure 2.9. Then, an additional term to describe the atom-laser interaction, $\hat{H}_{\text{int,L}}$, arises. So far, the total Hamiltonian of the

³⁴ A Raman transition is a two-step process: go from $|a\rangle$ to $|b\rangle$, then $|b\rangle$ to $|c\rangle$. Typically, this process is inelastic. However, the case we will investigate has energies $E_a = E_c$.

³⁵ Named for its shape resembling the character ‘ Λ ’.




Figure 2.9: Λ -level atom-cavity system. **(a)** A laser pulse with time-dependent Rabi frequency $\Omega(t)$ strikes the atom through the cavity's open side. **(b)** Energy levels of this system in the interaction picture (adapted from [55]). The laser and cavity couplings achieve a coherent state transfer from $|u, 0\rangle$ to $|g, 1\rangle$, during which the cavity emits a single photon at a rate 2κ . If we change $\Omega(t)$ adiabatically, the excited state $|e, 0\rangle$ is not populated throughout. Δ_L and Δ_C are laser and cavity detunings respectively.

system is (cf. equation 2.21)

$$\hat{H} = \hat{H}_a + \hat{H}_c + \hat{H}_{\text{int,L}} + \hat{H}_{\text{int,c}}, \quad (2.56)$$

where we have denoted the atom-cavity interaction \hat{H}_{int} (equation 2.29) from the previous analysis as $\hat{H}_{\text{int,c}}$. What follows now is to find the interaction Hamiltonian of the Λ -system so that we can reveal how this system can work as a single photon source.

We proceed with caution: defining zero energy for the atom as being halfway between $|g\rangle$ and $|e\rangle$, as we had done before, is potentially cumbersome when calculating the energy of $|u\rangle$.³⁶ Therefore, we instead define \hat{H}_a for the new system in terms of the absolute energy of the atomic levels:

$$\hat{H}_a = \hbar \left(\omega_u |u\rangle\langle u| + \omega_e |e\rangle\langle e| + \omega_g |g\rangle\langle g| \right). \quad (2.57)$$

The atom-laser interaction is straightforward: $\hat{H}_{\text{int,L}}$ describes the interaction of the atom

³⁶ A good example of calculating the interaction Hamiltonian for a Λ -level system in that case is shown in [56].

with a laser that strikes through the open side of the cavity, as illustrated in figure 2.9. Its fundamental difference to $\hat{H}_{\text{int,c}}$ is only that the light field is a free space laser mode, rather than the cavity mode. The derivations of $\hat{H}_{\text{int,c}}$ and $\hat{H}_{\text{int,L}}$ are therefore very similar, so we simply state (in the Schrödinger picture and with the RWA applied) [50]

$$\hat{H}_{\text{int,L}} = \hbar \frac{\Omega(t)}{2} \left(e^{-i\omega_L t} |e\rangle\langle u| + e^{i\omega_L t} |u\rangle\langle e| \right), \quad (2.58)$$

where $\Omega(t)$ is the time-dependent Rabi frequency of the laser pulse and ω_L is the laser frequency.

Again, we want to focus on the interaction now. Following the same process as in section 2.2.1, we move to the interaction picture. Since the atom can be driven off-resonantly by both the laser and the cavity, we accordingly define

$$\hat{H}_{\text{a},\Delta} = \hbar \left((\omega_u + \Delta_L) |u\rangle\langle u| + \omega_e |e\rangle\langle e| + (\omega_g + \Delta_C) |g\rangle\langle g| \right), \quad (2.59)$$

where $\Delta_L = \omega_e - \omega_u - \omega_L$ and $\Delta_C = \omega_e - \omega_g - \omega_c$. So, the associated interaction Hamiltonian is

$$\hat{H}_{\text{L},\Lambda} = \hat{U}_0^\dagger \left(\hat{H} - \hat{H}_{\text{a},\Delta} - \hat{H}_c \right) \hat{U}_0. \quad (2.60)$$

Completing the process in the same way as before, the resulting Hamiltonian in the interaction picture for $n = 1$ is, in the basis $\{|u, 0\rangle, |e, 0\rangle, |g, 1\rangle\}$ ³⁷ and with the RWA applied,

$$\hat{H}_{\text{L},\Lambda} = \hbar \begin{pmatrix} -\Delta_L & \Omega/2 & 0 \\ \Omega/2 & 0 & g \\ 0 & g & -\Delta_C \end{pmatrix}. \quad (2.61)$$

³⁷ Note the difference in ordering of the atomic states compared to the previous basis, $\{|g\rangle, |e\rangle\}$.

We have omitted the time dependence from $\Omega(t)$ for clarity.

The interaction similarly dresses the energy eigenstates, except now resulting in a triplet rather than a doublet. It is particularly instructive to calculate these eigenstates on two-photon (Raman) resonance, i. e. $\Delta = \Delta_C = \Delta_L$: they are (see appendix A)

$$\begin{aligned}
 |\psi^0\rangle &= \cos \Theta |u, 0\rangle + \sin \Theta |g, 1\rangle, \\
 |\psi^+\rangle &= \sin \Theta \cos \Phi |u, 0\rangle + \sin \Phi |e, 0\rangle - \cos \Theta \cos \Phi |g, 1\rangle, \\
 |\psi^-\rangle &= \sin \Theta \sin \Phi |u, 0\rangle - \cos \Phi |e, 0\rangle - \cos \Theta \sin \Phi |g, 1\rangle,
 \end{aligned} \tag{2.62}$$

where the mixing angles Θ and Φ are

$$\tan \Theta = \frac{\Omega}{-2g}, \quad \tan \Phi = \frac{-\sqrt{4g^2 + \Omega^2}}{\sqrt{4g^2 + \Omega^2 + \Delta^2 + \Delta}}. \tag{2.63}$$

The energy eigenfrequencies for this triplet of states for n quanta are

$$\begin{aligned}
 \omega_0 &= \omega_C \left(n + \frac{1}{2} \right), \\
 \omega_{\pm} &= \omega_C \left(n + \frac{1}{2} \right) + \frac{\Delta \mp \sqrt{4g^2 n + \Omega^2 + \Delta^2}}{2}.
 \end{aligned} \tag{2.64}$$

Two interesting subtleties arise. Firstly, $|\psi^0\rangle$ does not contain any term with $|e\rangle$. It therefore has no probability of the atom being excited: it is a dark state. Secondly, if $\Omega \rightarrow 0$, the system reduces to the previously found two-level Jaynes-Cummings system, as one would expect.

STIRAP

The presence of a dark state is interesting. Is there a way to transfer population from $|u, 0\rangle$ to $|g, 1\rangle$ whilst remaining in the dark state? Moreover, this would make emission of a single

photon possible, since $|g, 1\rangle \rightarrow |g, 0\rangle$ via the cavity decay.

Such a process is indeed possible. The ingredients for this coherent control process are stimulated emission, a Raman transition via $|e\rangle$, and adiabaticity. The recipe is therefore termed stimulated emission by Raman adiabatic passage (STIRAP). We will discuss the process qualitatively; a mathematical analysis reveals the precise conditions that make STIRAP possible (see reviews such as [57]), but we need only focus on the qualities necessary for our experiments.

The key insight of STIRAP is thus: it is possible to vary Θ —and therefore the relative population in either $|u, 0\rangle$ or $|g, 1\rangle$, as suggested by equation 2.62—and yet stay in the dark state throughout. The condition is that the dark state, $|\psi^0\rangle$, does not couple to either of the ‘bright’ states, $|\psi^\pm\rangle$. To fulfil this condition, Θ must be varied adiabatically. Adiabatic evolution corresponds to a slow, smooth change in $\Omega(t)$: a rough condition for adiabaticity is $\frac{d\Theta}{dt} \ll \Omega(t)$.³⁸ So, the process to do so is roughly as follows:

- ‘Switch on’ the interaction by adiabatically ramping up $\Omega(t)$ from zero.
- Change Θ by $\pi/2$ by controlling the ratio of the laser coupling to the cavity coupling: that is, increase $\Omega(t)$ to the required level.³⁹

Remarkably, since the state control remains coherent and the excited state is not populated throughout, loss due to spontaneous emission does not occur. There is no opportunity for population to build up in the excited state, since the cavity coupling effectively draws it back down to the ground state before it can build up. We may think of STIRAP as putting amplitude, but not population, in $|e\rangle$, which under adiabatic evolution will get drawn down to the ground state before any marginal population can build up (since population is amplitude

³⁸ This is the ‘local’ adiabatic condition: alternatively, the ‘global’ condition ensures that the pulse area (integral of Rabi frequency over the time of the pulse) is sufficiently large (see e. g. [58]).

³⁹ Afterwards, we may ‘switch off’ the interaction by adiabatically ramping $\Omega(t)$ back down to zero, but by then the photon has already been emitted, so doing so is not essential to complete the process.

squared).

During the **STIRAP** process, the cavity photon number increases by 1, since $|u, 0\rangle \rightarrow |g, 1\rangle$. Cavity decay by 2κ then brings $|g, 1\rangle \rightarrow |g, 0\rangle$, emitting the single photon from the cavity. In this case, since there are no photons initially in the cavity (referred to as vacuum-STIRAP (V-STIRAP)), we only emit the photon if we fire the laser pulse. Thus, we have created a deterministic single photon emitter. We can imagine pushing a button to emit a single photon on demand. Figure 2.10 shows this process when we apply a laser pulse $\Omega(t)$ with an envelope $\propto \sin^2$.

Under certain conditions, we may use the same **V-STIRAP** process to absorb/store a single photon [25, 59, 60]. If the cavity starts in $|g, 0\rangle$, then by coupling in a photon from, say, an optical fibre,⁴⁰ the laser pulse will similarly transfer the population from $|g, 1\rangle$ to $|u, 0\rangle$, resulting in the absorption of a single photon. This is the reverse of the photon generation process. So in addition, we have created a deterministic single photon absorber. Usually, we require impedance matching (eliminate any field reflections as the photon impinges on the cavity by suitably driving the atom with the laser pulse) to do this effectively. Nevertheless, since it is strong cavity coupling which makes the process efficient, then any cavity that efficiently generates photons should in principle be able to efficiently absorb photons too [61].

With respect to the detunings, we need two-photon (Raman) resonance $\Delta_C = \Delta_L$ to ensure that the dark eigenstate does not couple to any other [58], although we can tolerate small deviations from this.⁴¹ The Φ mixing angle quantifies the effect of the detunings: we can use this to strike a balance between reducing the chance of a radiative process and increasing the effective coupling between the desired initial and final states of the **STIRAP** process. Therefore, the detunings are crucial to an efficient process. We will examine this numerically

⁴⁰ At a rate κ_{in} , as according to the input-output formalism that we will discuss in chapter 3.

⁴¹ When not Raman resonant, we may think of the system as precessing around the dark eigenstate; the precession gets stronger as we drift further from Raman resonance.

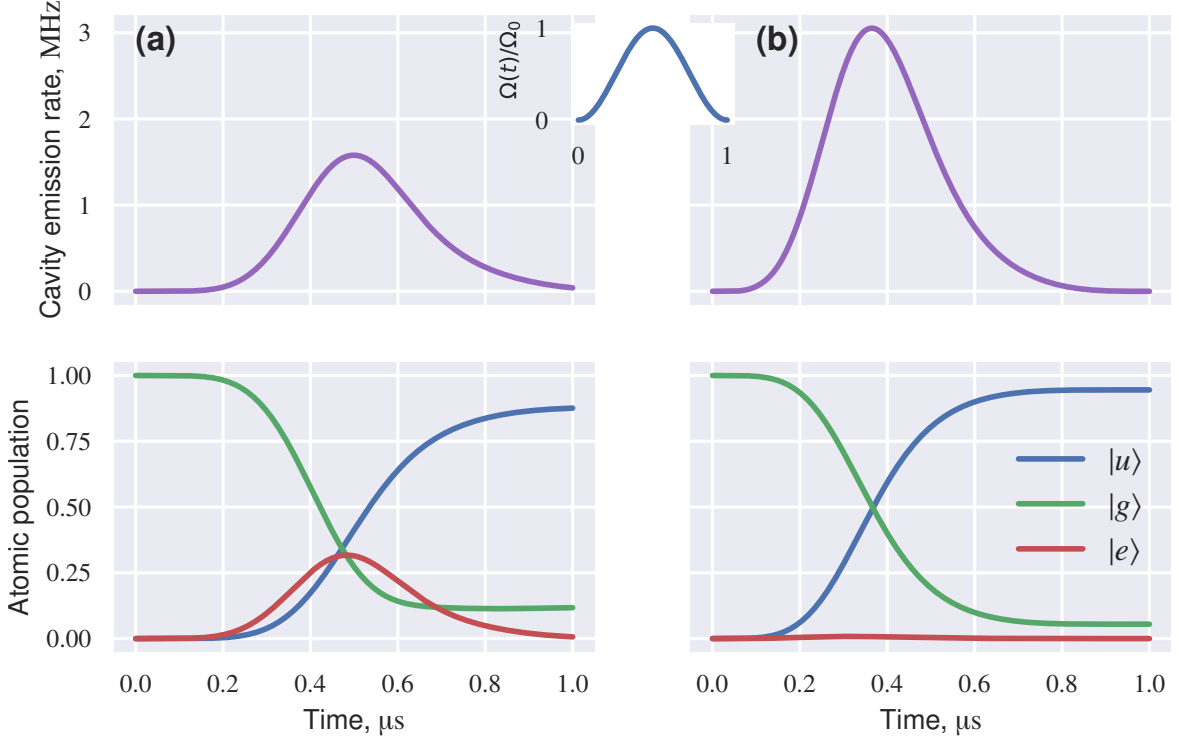


Figure 2.10: Atom-cavity system dynamics for single photon emission during the **STIRAP** process, starting from $|u, 0\rangle$. **(a)** Bad cavity regime, with $(g, \Omega_0, \kappa, \gamma)/2\pi = (2, 2, 10, 1)$ MHz. **(b)** Strong coupling regime, with $(g, \Omega_0, \kappa, \gamma)/2\pi = (10, 10, 3, 3)$ MHz. We apply a laser pulse with a \sin^2 envelope (inset), and no detunings (to better illustrate the process). To plot these figures, we numerically solve a master equation, as we will describe in chapter 3.

in chapter 3; it has also been examined in the the literature [57, 62].

With respect to the couplings, increasing g and/or decreasing κ and γ will only serve to increase the efficiency of the photon generation process. From equation 2.55, we see that increasing the co-operativity C of this system gives the best possible single photon source and memory, as discussed prior.⁴² Although the **STIRAP** process is quite robust to small changes in the laser power or detunings, we must also ensure the peak power of the laser pulse is large enough to overcome the incoherent processes that typically occur in experiments over

⁴² Following figure 2.10b, we also see that the emitted photon shape is asymmetric. This is because we must drive the system harder at later times to retrieve the ever-decreasing population at a constant rate, akin to squeezing a sponge harder to extract the last drops of water. Therefore, a symmetric $\Omega(t)$ yields an asymmetrically shaped photon. We will investigate the effects of this asymmetry on our source/memory performance in chapter 3.

a longer timescale. Effectively, this means we must keep the pulse length shorter than the coherence time of the system.

Altogether, in the strong coupling regime and with the optimal parameters, the probability of deterministically emitting a single photon approaches unity. Chapter 3 will explore how efficiently we can generate single photons with the cavities in our experiments.

2.3 Summary

We have defined the framework for the experiments in the following chapters: that framework is [CQED](#). From first examining how an optical cavity can be used to magnify and concentrate the intensity and frequency of light to realise a very strong field for a single mode, we have found the form of the field in the cavity. We then placed an atom in the cavity and quantised the field, from which we have determined the nature of the atom-light interaction in the cavity at the single quantum level. Remarkably, the states of the atom and light couple to each other, giving a constant exchange of a single excitation between them. We examined how this exchange is effected in realistic experiments: both when we changed the parameters of the coupling (e. g. off-resonant driving) and when we introduced the effect of system losses to the environment. From there, we demonstrated that when the coupling is strong, the exchange almost exclusively channels the resulting photon emission into the cavity mode rather than the environment. Lastly, we combined all of these properties with a [V-STIRAP](#) process, showing that the system can effectively realise a single-photon source, and what further conditions are necessary to realise this. Thus, we demonstrate how atom-photon interfacing can be realised in the simulations and experiments in the following chapters.

Optimal parameters for cavity-based quantum networks of atoms and photons

Note: this chapter has been prepared in the form of a journal paper, minus section 3.4, which remains in this thesis to provide important context for the other results presented here.

Abstract

To realise quantum networks of atoms and photons, one may use cavity-based atom-photon interfacing to share quantum information in a faithful, controllable manner. However, the limits of such interfacing depends on the parameters of the cavities and laser pulses used. Do such limits fundamentally restrict the capability of cavity-based quantum networks? Here, we examine this by numerically optimising the parameters of systems of cavities each coupled to a single atom. We optimise for generating and absorbing single photons with maximal efficiency and minimal losses. Firstly, we optimise the atom-cavity system parameters, and secondly the shape of the laser pulse, both using the Nelder-Mead algorithm. For a [fibre Fabry-Pérot cavity \(FFPC\)](#) with realistic parameters, we achieve 99.1 % generation efficiency with 98.1 % coherence for photon lengths ranging from 1 to 0.3 μs . For absorption, we achieve 62.8 % efficiency with 99.6 % coherence. Thirdly, we optimise remote state transfer from one cavity to another by a single photon, achieving 39 % transfer efficiency for highly directional

cavities (one mirror more reflective than the other), and 20 % for the opposite case, gaining crucial insights into how to find the optimum directionality for a given network. We also explore the application of the Krotov algorithm to the pulse shape optimisation, elucidating how irreversible cavity emission hinders the application of gradient-based algorithms to these systems. Overall, our optimisation will aid in realising efficient quantum networks with well-controlled and coherent transfers.

3.1 Introduction

Quantum networks, which can connect several quantum processors via quantum channels to harness their collective capability [14, 21, 33], can potentially allow us to study larger superposed and entangled states [63, 64]. Networks at the scale of a few nodes have been realised in the laboratory [23, 65]. In Ritter et al. [23], optical cavities constitute atom-photon interfaces and are all linked by optical fibres, as illustrated in figure 1.1. There, a logical qubit is either the electronic state of an atom ($|M_F = \pm 1\rangle$ hyperfine sublevels)¹ or the polarisation state of a cavity photon (left(right)-handed circular polarisations driving $\sigma^+(\sigma^-)$ transitions with $\Delta M_F = +1(-1)$). Prior to Ritter et al.’s research, Kimble had questioned ‘whether a quantum network “works”’ [14], summarising today how critical it is to faithfully control the states across the network as it continues to grow in size.

Such quantum networks are not limited to neutral atoms in cavities, but they must make use of stationary and travelling qubits for information storage and communication respectively. Photons are the natural candidates for the latter since they can maintain their coherence over longer distances [15], but the candidate for the former is a more open question. Well-suited types of stationary qubits include, but are not limited to: ions, either with cavities [66] or

¹ Where M_F is the z -component of the total angular momentum of the atom (choosing the cavity axis as the z -axis).

without [67]; color centres, again with cavities [68] or without [65]; quantum dots [69] or other solid state systems [70]; and atomic ensembles [71]. As is evident from these references, cavities are not required for quantum networks. At a high level, the cavity's main purpose is to 'trap' the travelling photon and enhance its interaction with the stationary qubit, but especially with respect to atomic ensembles, new, efficient schemes have been performed in the laboratory that suggest significant thresholds of performance may be fundamentally attainable without cavities [72]. But, atoms are still the fundamental building blocks of all of these systems. So, we may expect that all these network types share similar limitations. Our pertinent question then becomes: can we use atoms in cavities to investigate the fundamental limits relevant to quantum networks in general?

To this end, in this simulation work we optimise atom-cavity systems for cavity-based quantum networks. Defining the optimisation goal, we must optimise the photon generation/absorption efficiencies, losses, photon lengths and coherences for the systems' parameters.

Firstly, we optimise the parameters of a single system on its own. Although the parameters' complex interplay can make an analytical optimisation non-trivial, fortunately the total number of system variables is small ($N \lesssim 20$) and therefore we can efficiently apply numerical optimisation methods. We choose the Nelder-Mead algorithm [73, 74]: a standard method that performs well when applied to 'black box' problems such as this, where we assume no knowledge of how to find the optimal parameters.

Secondly, we optimise the shape/envelope of the laser pulse that drives the atom-photon interfacing process. In particular, the laser drives a **STIRAP** process between two ground states of the atom which adds (or removes) a single photon to (or from) the cavity. We stress that previous analytic optimisations of the pulse shape do not optimise the process's efficiency or coherence. Instead, they optimise a related quality, such as finding what pulse shape will yield: a certain photon shape [75]; impedance matching at the cavity input/output [59]; or

non-population of the excited state [76]. All rely on simplifying assumptions to achieve their result, and so they are not necessarily the most efficient shapes to do so. Our numerical approach circumvents these assumptions by encapsulating the desired photon properties and process outcome in a single function of performance. Note that we optimise both the system parameters and pulse shapes, rather than just one or the other as has been performed before for experiments [77, 78]. It also grants us the novelty of optimising networks containing non-identical nodes: for example, hybrid networks that mix and match nodes of multiple types of atom-cavity systems (or not with cavities [79, 80]). Realistically, every cavity network contains non-identical nodes because of manufacturing variations, e.g. mirror coatings, mirror geometric shapes.

As well as applying the Nelder-Mead algorithm, we also investigate and apply the Krotov algorithm [81, 82] to pulse shape optimisation. We find that typical gradient-based algorithms, such as gradient ascent pulse engineering (GRAPE) and Krotov, may not be suitable to optimising our particular **STIRAP** process, because of our reliance on dissipation to emit the photon from the cavity.

The structure of this paper is as follows. We begin by describing the ‘building block’ elements of our simulations in section 3.2. Then, we move onto the optimisation. Firstly, we optimise the parameters of a single atom-cavity system for photon generation and absorption in section 3.3. Secondly, we optimise the pulse shape using Krotov’s algorithm in section 3.4. Thirdly, we optimise the pulse shape using the Nelder-Mead algorithm in section 3.5. Fourthly, in section 3.6 we optimise remote state transfer where a photon emitted from one cavity is absorbed in another. Finally, we conclude in section 3.7.

3.2 Simulation elements

Consider the Λ -level atom-cavity system that we described in figure 2.9 and section 2.2.3, by which we use **STIRAP** to emit(absorb) a single photon into(from) the cavity. To find the optimal atom-cavity parameters, we must simulate the **STIRAP** process. This necessitates calculating how the system's state varies in time. We must expand on our previous model by including dissipation, especially since our **STIRAP** process relies on cavity decay. To do this, we use the master equation in Lindblad form [83],

$$\frac{d\hat{\rho}}{dt} = -i [\hat{H}, \hat{\rho}(t)] + \hat{D}(\hat{\rho}(t)), \quad (3.1)$$

where $\hbar = 1$,² \hat{H} is the system Hamiltonian (cf. equation 2.61), $\hat{\rho}(t)$ is the system's density matrix at time t , and \hat{D} is the dissipation superoperator that includes the decays through collapse operators \hat{C}_n :

$$\hat{D}(\hat{\rho}(t)) = \frac{1}{2} \sum_n \left(2\hat{C}_n \hat{\rho}(t) \hat{C}_n^\dagger - \hat{C}_n^\dagger \hat{C}_n \hat{\rho}(t) - \hat{\rho}(t) \hat{C}_n^\dagger \hat{C}_n \right). \quad (3.2)$$

Each \hat{C}_n represents one path of decay. For example, one path could be spontaneous emission by the atom from excited state $|e\rangle$ to ground state $|u\rangle$, which is $\sqrt{2\gamma}|u\rangle\langle e|$.

If we model photon absorption rather than generation, then we also need to add a corresponding coupling for the absorption. A single photon impinging on the cavity mirror couples in to the cavity, and the **STIRAP** laser pulse subsequently causes the atom to absorb the photon. So, we add an impinging photon in-coupling to the Hamiltonian, where the coupling is of the form [61, 84]

$$\hat{\phi}_{\text{in}}(t) = -\mathcal{N}S(t)\sqrt{2\kappa_{\text{in}}}\left(\hat{c} + \hat{c}^\dagger\right), \quad (3.3)$$

² We will set $\hbar = 1$ in this chapter only.

where \mathcal{N} is a normalisation factor such that for shape (i.e. intensity envelope) S , the impinging wavepacket contains exactly one photon, and where $2\kappa_{\text{in}}$ is the (intensity) incoupling rate through the mirror that the photon impinges on. Let us expand on how we implement these two aspects.

In order to normalise the impinging single-photon wavepacket with probability amplitude $\psi(t)$, we must have $\int_0^T |\psi(t)|^2 dt = 1$, where T is the length of our simulation. Therefore, its shape $S(t)$ —its intensity envelope—should satisfy

$$\int_0^T S(t) dt = A \int_0^T |\psi(t)|^2 dt, \quad (3.4)$$

since intensity is proportional to amplitude squared, and where A is a proportionality constant. Since we are considering a single photon ($N = 1$), we also have

$$\mathcal{N} \int_0^T S(t) dt = \sqrt{N} = 1. \quad (3.5)$$

Combining equations 3.4 and 3.5, we find $\mathcal{N} = A$, which is as we wanted to demonstrate: normalising $S(t)$ by \mathcal{N} ensures the wavepacket contains exactly one photon.

To determine the incoupling rate $2\kappa_{\text{in}}$, we have used the input-output formalism as in [61, 84]. Usually, this involves integrating a frequency-dependent coupling parameter over all possible modes. However, our high-finesse cavity has a correspondingly low density of photon modes—it effectively ‘picks out’ a frequency of interest. Thus, we may make the simplification that the impinging photon only couples to the single cavity-resonant mode. After doing this, we will investigate two possible cases:

- *Directional cavity*: one cavity mirror is much more reflective than the other. This leaves only one mirror decay channel, thus $\kappa_{\text{in}} \approx \kappa$, where 2κ is the cavity linewidth.

- *Symmetric cavity*: both cavity mirrors have the same reflectivity. Here, $\kappa_{\text{in}} \neq \kappa$.

We consider the first case because it matches the model by Cirac et al. [25] when the mirror losses are small and when connecting two cavities together with a fibre, thus we can more readily compare our results with those. However, it is not always realistic or possible to ignore the small mirror losses or the decay channel through the second mirror, which is why we consider the second case. To find κ_{in} in that second case, we consider the ratio of how much light couples into the cavity versus how much couples in, couples out and is lost:

$$\frac{\kappa_{\text{in}}}{\kappa} = \frac{T_{\text{L}}}{T_{\text{L}} + T_{\text{R}} + \mathcal{L}}, \quad (3.6)$$

where T and \mathcal{L} are the mirror transmission (intensity) and loss respectively, and where the L and R subscripts denote the left and right mirrors. Because we want to investigate how a realistic cavity would fare in a cavity-based quantum network, we therefore use the parameters of a [fibre Fabry-Pérot cavity \(FFPC\)](#) that we built and implemented in the laboratory. Calculating κ_{in} for our fibre cavity using the parameters we measured in [chapter 4](#) and given in [table 4.3](#) gives $\kappa_{\text{in}} = 0.35\kappa$.

Following the above, throughout this work we will use the first case of $\kappa_{\text{in}} = \kappa$, except for in [section 3.6](#) where we will consider both the first and second cases in order to examine what effect using a realistic non-directional cavity has on the optimal networking parameters.

To simulate the process, we need to solve [equation 3.1](#) from time 0 to a final time T . We do this using a modified version of the `rb-cqed` Python package [85], which creates the Hamiltonian based on what energy levels, couplings etc. we set, and then passes it to `qutip.mesolve` in the QuTiP Python package to numerically solve the master equation. Compared to the interaction picture Hamiltonian we defined in [equation 2.61](#), we use a rotating frame in which we have zero energy in the excited state and which is rotating at the

same frequency as the detunings. As one example, we replace $\hat{H}_{a,\Delta}$ with \hat{H}_a in equation 2.60. Thus, the atomic transitions to other levels will have small frequencies, making the numerical integration computationally easier. More information on the simulation back-end is also available in [32].

Note that we are aiming to optimise these simulated processes for quantum networks, which necessitates that we maintain the coherence. For this reason, we must carefully consider the effect of spontaneous atomic decay. If a spontaneous decay occurs from the excited state and the system is thereafter back in the initial state, the **STIRAP** process can restart, but incoherently. Therefore, in all of our simulations except in figure 3.1 at the beginning of section 3.3, we include a virtual ‘sink’ state to which all spontaneous emission is directed and which is decoupled from the system.

3.3 Optimal atom-cavity parameters

Firstly, we focus on photon generation. Our goal is to find the parameters for the **STIRAP** process that maximises generation efficiency of a single photon while minimising unwanted spontaneous emission. We encapsulate this goal in a performance function,

$$J = 1 - \eta + \Upsilon, \tag{3.7}$$

where η is the photon generation efficiency and Υ is the loss due to spontaneous emission, calculated by integrating the expectation values of the cavity photon number operator and spontaneous emission operators respectively from time 0 to T . We must minimise J .

Then, we use the Nelder-Mead algorithm [73, 74], to find the system parameters that minimise J . The algorithm works as follows: starting from the initial g , κ , Ω_0 , Δ_L and Δ_C (all defined in section 2.2.3; Ω_0 will be defined soon in equation 3.8), we construct a

volume formed of $N + 1 = 6$ points, known as a simplex, where each point has co-ordinates corresponding to a slight variation of one of the parameters. Next, we initialise the algorithm by evaluating J at each point to see which variation gives the minimal J . The task of the algorithm is then to iteratively replace the worst performing point (maximal J) by a new, better performing point, until all the points lie in a valley of minimal J . For one iteration, we reflect the worst point about the simplex's centre point, and evaluate J at the new point. Then, depending on how much better the point is, along the line connecting the two we either: expand the simplex; stick with the reflection; or contract the simplex towards the centre. If no improvement is made, we shrink the whole simplex towards the best performing point. Although, shrinkages are rare in practice [73]. Thus, the optimisation process is as follows:

1. Starting from the initial g , κ , Ω_0 , Δ_L and Δ_C , simulate the STIRAP process.
2. Evaluate $J(t = T)$ at the end of the simulation.
3. Perform one iteration of the Nelder-Mead algorithm to replace the worst performing point with a better point(s), running a simulation again once (or sometimes twice, or rarely N times) to do this.
4. Repeat steps 2 and 3 to minimise J .

We terminate the optimisation after 200 iterations, because it converges to the optimal solution after about 100 iterations. To put this in perspective, for an n -parameter problem we would very roughly expect to achieve convergence after $100n$ iterations, all with typically with only one or sometimes ≥ 2 simulations per iteration [86]. Thus, our advantage of using Nelder-Mead here is the ability to achieve a decent optimisation without having to run the computationally expensive simulation many times over.

Additionally, we bound parameters to realistic values for a ^{87}Rb atom in a fibre-tip Fabry-Pérot cavity, in line with what we achieve for our cavity assembly in chapter 4 and table 4.3.

Parameter	Initial value	Bounds	Optimal value
g	10	[5, 50]	50
κ	4	[3.5, 100]	6.8
Ω_0	30	[25, 50]	35
Δ_L, Δ_C	0	[-50, 50]	6.0, 7.1

Table 3.1: Initial and optimised values, and bounds, for the atom-cavity parameters. All values are in units of $2\pi \times \text{MHz}$. The atomic spontaneous field decay rate $\gamma/2\pi$ is fixed at 3.0 MHz.

The bounds are in table 3.1. Using the `constrNMPy` Python package [87], at the start of an iteration we internally transform each parameter to keep them within the bounds; then at the end of an iteration, we perform the inverse transform to retrieve them. Note that our faster than expected convergence may be due to using these bounds: specifically, bounding reduces the parameter space that the algorithm would explore.

The cavity coupling is on the $^{87}\text{Rb D}_2$ line; here, the atomic states $|u\rangle$, $|g\rangle$ and $|e\rangle$, as defined in section 2.2.3, correspond to the hyperfine ground states $|F, m_F\rangle = |2, -1\rangle$, $|1, +1\rangle$, and excited state $|F', m_{F'}\rangle = |1, 0\rangle$ respectively.³ The laser coupling is of the form

$$\Omega(t) = \Omega_0 S(t), \quad (3.8)$$

where Ω_0 is the peak laser Rabi frequency and $S(t)$ is the envelope or shape of the pulse. Intuitively, we may also think of Ω_0 as corresponding to the laser power. In this section, we keep the shape fixed to $S(t) = \sin^2(\pi t/T)$, which we choose as a good initial guess for a **STIRAP** process because it is a smooth shape without sharp gradient changes and which tends to zero value and zero gradient at the pulse extremities.

Now, we perform the optimisation of a single photon generation process. Starting from the system state $|u, 0\rangle$, we aim for the system to reach $|g, 0\rangle$, after which a single photon will have been emitted from the cavity. We choose the initial parameters in table 3.1 as a good guess

³ For brevity here, we defer more details about hyperfine states and the $^{87}\text{Rb D}_2$ line to section 5.2.4.

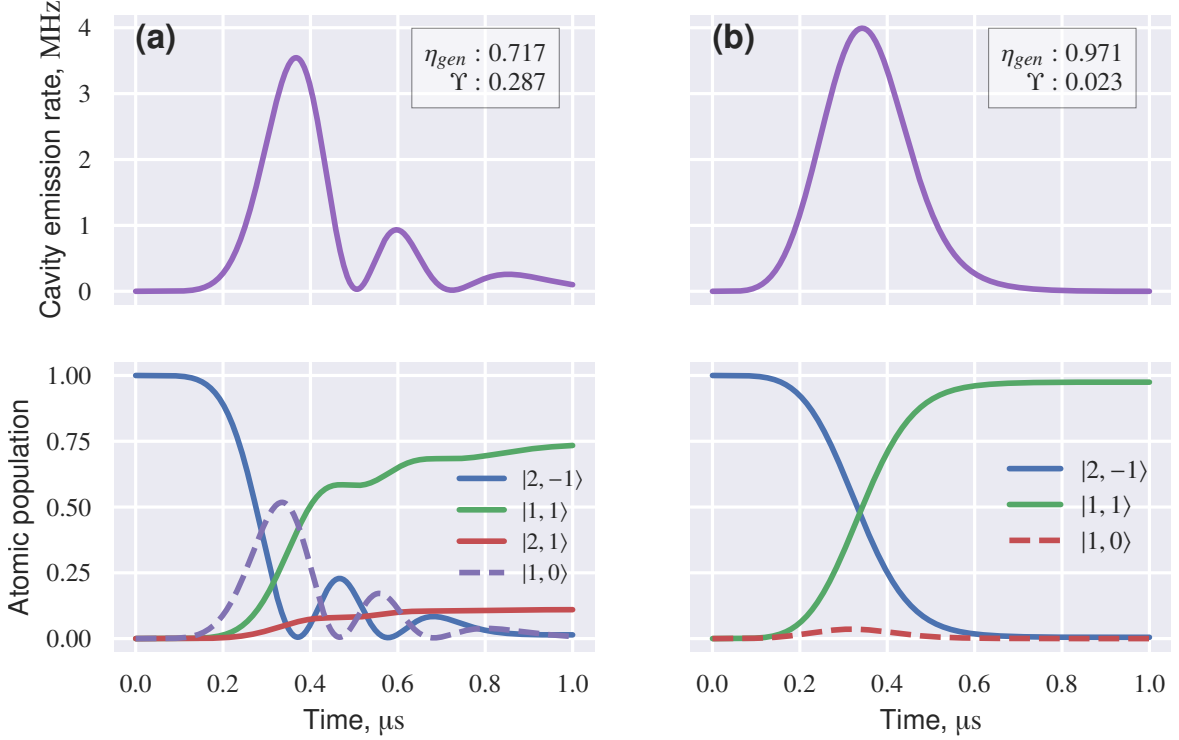


Figure 3.1: Simulated photon generations for: **(a)** the initial parameters; and **(b)** the optimal parameters (both in table 3.1). The laser pulse couples $|u\rangle$ and $|e\rangle$, and the cavity couples $|e\rangle$ and $|g\rangle$. We see the shape of the emitted photon in the upper plots. Here, we simulate the full hyperfine structure of the ^{87}Rb D_2 line including the Zeeman sublevels, totalling 24 sublevels [88], although for clarity we omit most of the levels from the plot as they do not make a significant contribution. We denote atomic ground states $|F, M_F\rangle$ with solid lines and excited states $|F', M'_F\rangle$ with dashed lines. As mentioned in section 3.2, we do not use a sink state for these particular simulations.

for an efficient process, being in the strong coupling regime $g > \kappa, \gamma$. Figure 3.1a shows the initial dynamics. Due to off-resonant pumping, some population gets stuck in an undesired ground state $|2, 1\rangle$.

After completing the optimisation, we have the final dynamics in figure 3.1b. We see that the **STIRAP** generation process works as intended, as the system goes from $|u, 0\rangle$ to $|g, 0\rangle$ with the emission of a photon and without populating $|e, 0\rangle$. We have improved η from 71.7 % to 97.1 % and Υ from 28.3 % to 2.3 %. Note here that we do not use a sink state, since we are interested in seeing which atomic levels and decay paths are the dominant ones in our realistic atom with many couplings. Without the sink state, in figure 3.1a we see a strong cavity

coupling from multiple Zeeman substates (via $\Delta M_F = \pm 1$ transitions) after a spontaneous decay. Especially in this case where the STIRAP process is only partially successful, we must therefore be aware that a strong cavity coupling can work against us when preparing an atomic in a particular Zeeman state (e.g. superposition of $|2, \pm 1\rangle$, as is used for the logical qubit).

When running the optimisation again with a sink state, we now see that we maximise η from 12 % to 95.4 % and minimise \mathcal{Y} from 10 % to 6.6 %. Tracking the convergence of the individual parameters, we obtain figure 3.2, from which we see that:

- g reaches its upper bound (since the process should always be improved by stronger cavity coupling);
- κ reaches an optimal balance between cavity coupling and photon emission;
- Δ_L and Δ_C are kept close to Raman resonance throughout;
- Ω_0 reaches an optimal balance between strongly driving the system and not undesirably fluorescing/pumping the atom.

With respect to g and κ , maximising g and allowing the photon to escape with a reasonable κ is what we may intuitively expect for an efficient process [39]. With respect to Ω_0 , we have effectively penalised unwanted fluorescence using \mathcal{Y} , because when pumping occurs, \mathcal{Y} increases more than η does, and more than one photon would be incoherently emitted. Thus, we especially see the importance of Ω_0 to the efficiency.

We note also that our optimal parameters are close to, but not exactly at, Raman resonance $\Delta_L = \Delta_C$. Usually, Raman resonance is a strict condition for STIRAP [58]. To investigate this, we scan the detunings from the optimal parameters, as shown in figure 3.3. Our result matches that of Kuhn et al. [62]. We see that there is a small amount of leeway in the detunings while still maintaining good performance. Our observation is supported by the fact

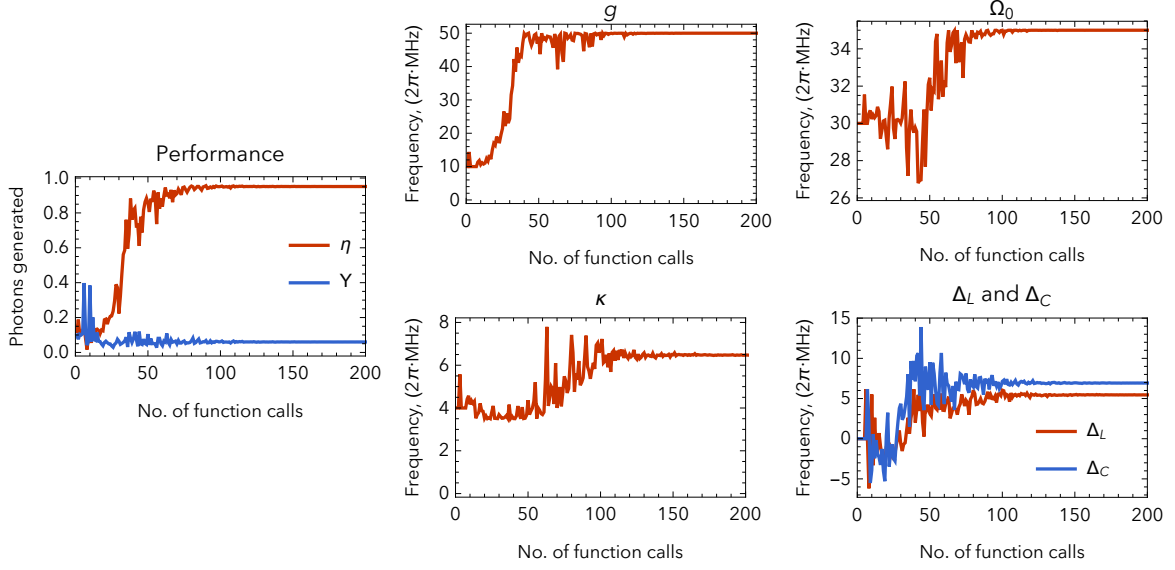


Figure 3.2: Convergences of the optimisation performance and of the atom-cavity parameters. Here, we direct all spontaneous emission to a virtual sink state that is decoupled from the system (see main text). We plot versus the number of function calls rather than iterations because a single iteration may have one or many function calls, the latter of which could mask a potentially large jump in the parameters.

that an atomic resonance has a natural linewidth 2γ and is not perfectly sharp, and so we can tolerate small detunings well below γ . We have more leeway than in [62] because our system parameters are further into the strong coupling regime.

We also simulate photon absorption using the initial parameters and optimised parameters from table 3.1, as shown in figure 3.4. The efficiencies and spontaneous emission are lower than for generation, with the photon absorption efficiency replacing the generation efficiency, i.e. $\eta_{\text{abs}} \approx \eta_{\text{gen}}$. We improve the initial η_{abs} from 13.8% to 39.4%. One explanation is that we are not explicitly setting impedance matching: as we see in figure 3.4, significant cavity emission still takes place during the optimal absorption. We note also that η_{abs} is more sensitive to the two-photon detuning: for this plot, we have set $\Delta_L/2\pi = \Delta_C/2\pi = 6.0$ MHz, noticing a slight boost in efficiency of about 2%. Considering these observations, we expect that a cavity well-suited to single photon generation may not be as well-suited to absorption

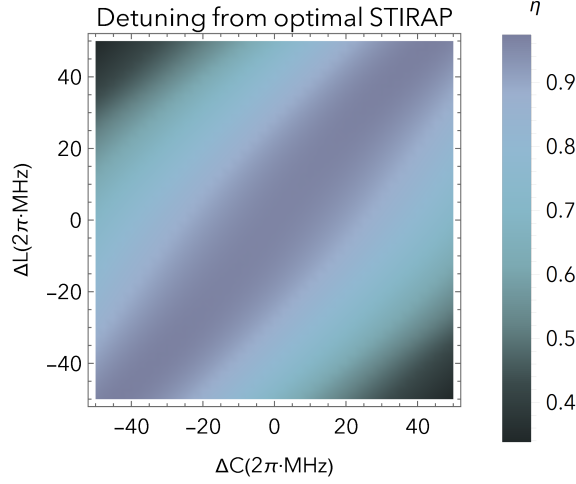


Figure 3.3: How the efficiency of photon generation η varies as we scan Δ_L and Δ_C from the optimal atom-cavity parameters (which are in table 3.1).

unless either the laser pulse is significantly different or the impedance matching condition is met. We will investigate the former when we optimise the pulse shape in the latter sections of this chapter. Although our many-level simulation incorporating the full atomic hyperfine structure does not yet facilitate photon absorption, by using our result and comparing it with the findings by Cirac et al. [25] and Dilley et al. [59], we find that the optimal parameters for absorption are still similar to those for generation (apart from the change to exact Raman resonance). One may expect this from time reversal symmetry since, qualitatively, the emitted and absorbed photon shapes for the optimal system parameters are fairly symmetric.

We emphasise that for the following sections we will not simulate the full hyperfine structure. This is because simulating the full structure is computationally taxing, and we now have the optimal system parameters: we only optimise pulse shapes from here on. We only simulate the 3 atomic levels relevant to **STIRAP**: $|u\rangle$, $|e\rangle$, and $|g\rangle$, as we defined in sections 3.2 and 2.2.3, as well as the previously mentioned sink level, which we will denote $|s\rangle$. Because the optimised dynamics are now efficient enough that off-resonant coupling to unwanted levels is low, we justify reducing the number of simulated levels.

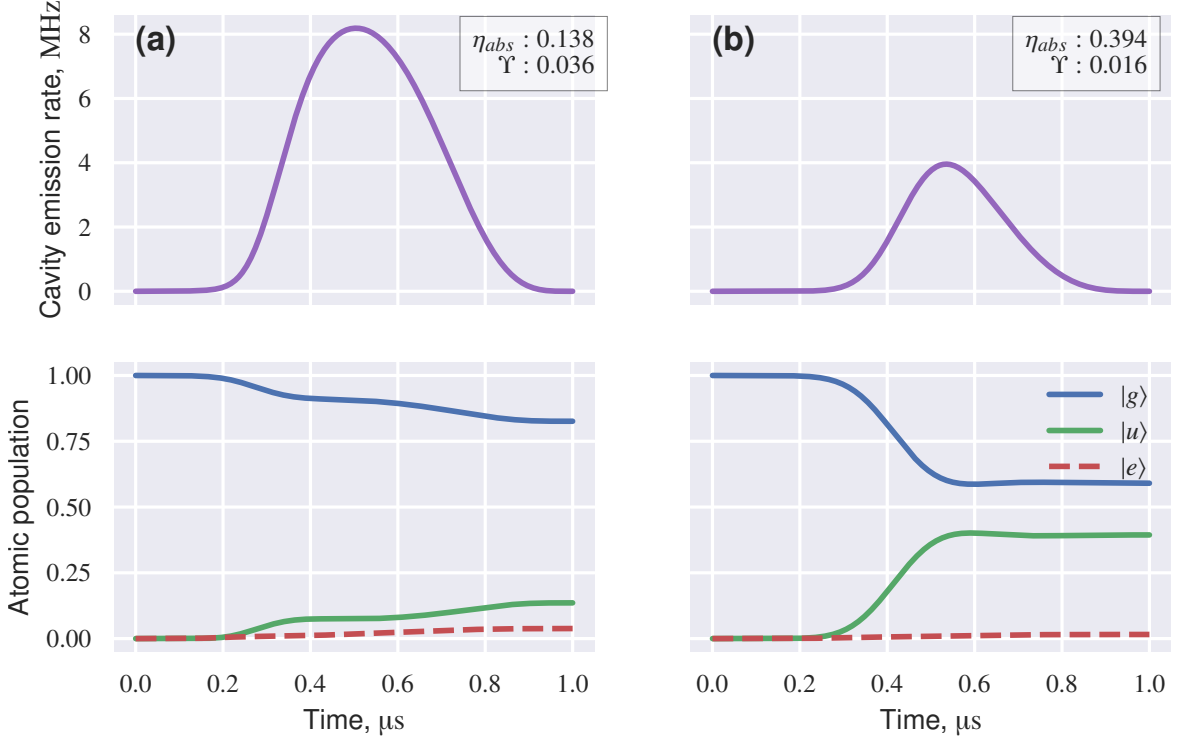


Figure 3.4: Simulated photon absorptions for: **(a)** the initial parameters; and **(b)** the optimal parameters (both in table 3.1), apart from setting Raman resonance $\Delta_L/2\pi = \Delta_C/2\pi = 6.0$ MHz. We do not simulate the full hyperfine structure here (see main text).

3.4 Krotov pulse shape optimisation

Consider the evolution of the atom-cavity system from $|u, 0\rangle$ to $|g, 0\rangle$. Since the pulse S steers the system along this trajectory, then the optimal pulse S_{opt} for generating a photon is the one that steers the system into $|g, 0\rangle$ with the highest fidelity. When framed in this way, we can use algorithms from quantum optimal control to achieve this goal [89]. These algorithms all arise from applying the variational principle to derive how to change the pulse to maximise this fidelity [90]. Algorithms such as GRAPE [91] and Krotov [81, 82, 92] are based on gradient descent: one can derive an analytical expression for the gradient of the performance function J with respect to the pulse, $\frac{\partial J}{\partial S}$, and therefore can ‘roll’ the pulse down this gradient to a minimum J (‘bottom of the valley’) and therefore maximum fidelity. Here, we apply the

Krotov algorithm to find the optimal S for our **STIRAP** processes.

In our application of Krotov, we construct S_{opt} by iteratively steering the trajectory towards the target state $|\psi_{\text{tgt}}\rangle$. Firstly, we propagate $|\psi(0)\rangle$ to $|\psi(T)\rangle$ under a good guess for the initial pulse. We term this a forward trajectory. Secondly, we project $|\psi(T)\rangle$ onto our target state $|\psi_{\text{tgt}}\rangle$. After that, we run a time-reversed evolution from there. We term this an optimising trajectory. By comparing the forward and the optimising trajectories, and operating on them both with an expression that gives us $\frac{\partial J}{\partial S}$ for each of them, we effectively modify S such that it steers the forward trajectory towards the target. We apply this iteratively, repeatedly making forward and optimising trajectories. Using this gradient descent to adjust $S(t)$ at every point in time, eventually the forward trajectory will match the optimising trajectory (apart from dissipation). At this point, the optimisation has converged/we have reached S_{opt} , since it steers into the target.

3.4.1 Setting up the Krotov algorithm

For the following formal definition/application of Krotov, we focus only on the salient aspects for our particular system: full details of the algorithm and the code that we adapt for our use are available in [81], from which we use the `krotov` Python package. For brevity and ease of understanding, we will reintroduce dissipation (i.e. collapse operators) later, only after we set up the algorithm. The reader may also wish to skip ahead to section 3.4.2 for an algorithm summary and a pictorial representation of it in figure 3.5.

For photon generation, our goal is to emit a single coherent photon by starting from

$$|\psi_{\text{init}}(t = 0)\rangle = |u, 0\rangle$$

and reaching

$$|\psi_{\text{tgt}}(t = T)\rangle = |g, 0\rangle$$

subject to the laser pulse that causes a coupling Ω as defined in equation 3.23.

Our system is governed by the Hamiltonian, which we can separate into a time-independent and a time-dependent part:

$$\hat{H}(t) = \hat{H}_0 + S(t)\Pi(t)\hat{H}_1, \quad (3.9)$$

where both H_0 and H_1 are time-independent. Because STIRAP requires a smooth pulse, we have incorporated a window function $\Pi(t)$ that smooths the pulse at the endpoints, ensuring that $\Omega = \frac{d\Omega}{dt} = 0$ at times 0 and T .⁴ Note that we have absorbed Ω_0 into \hat{H}_1 .

We shall now define the performance function in two steps. Firstly, its part at the end of the simulation is

$$J_T = 1 - F(T), \quad (3.10)$$

where the fidelity is

$$F(T) = \left| \langle \psi_{\text{tgt}}(T) | \psi(T) \rangle \right|^2. \quad (3.11)$$

So, we will reach S_{opt} when J_T is minimised. Secondly, we include a cost function, so that our full performance function is now

$$J = J_T + \int_0^T \mathbb{C}(S(t)) dt, \quad (3.12)$$

where \mathbb{C} is the cost or penalty incurred when changing $S(t)$,

$$\mathbb{C} \propto \Delta S(t)^2, \quad (3.13)$$

⁴ Here, $\Pi(t) \in [0, 1]$ is a flat-top with a half-sin² shape at the extremities, and rise and fall times of 50 ns.

and where $\Delta S(t)$ is the change in the pulse amplitude from $S(t)$. Thus, \mathbb{C} ensures that the ΔS made is the most efficient one possible: that is, the effort expended in changing S exclusively results in minimising J .⁵ We reinforce our picture so far by noting that the above calculates J for a single simulation: therefore, we should expect that we will apply the algorithm iteratively, repeatedly simulating the process and updating the pulse by $\Delta S(t)$ at every point in time once per simulation.

The optimal control problem is now clear: we must make the least costly change to $S(t)$ at every point in time $0 \leq t \leq T$ in order to minimise J .

For Krotov, we construct an ‘optimising’ trajectory to do this. In comparison to the forward trajectory $|\psi_{\text{init}}(0)\rangle \rightarrow |\psi(T)\rangle$, the optimising trajectory follows $\langle\chi(T)| \rightarrow \langle\chi(0)|$, where we choose the new state $\langle\chi(T)|$ by projecting the final state $|\psi(T)\rangle$ onto our target:

$$\langle\chi(T)| = -\langle\psi(T)|\psi_{\text{tgt}}\rangle\psi_{\text{tgt}}|. \quad (3.14)$$

By setting this boundary condition for the optimising trajectory, we can calculate the necessary pulse update ΔS that, at every point in time, will always follow the gradient $\frac{\partial J}{\partial S}$ and therefore improve J (monotonic convergence) using [81, 82, 92]:

$$\hat{\mu}(t) = \frac{\partial \hat{H}(t)}{\partial S} = \hat{H}_1, \quad (3.15)$$

$$\nu(t) = \text{Im}(\langle\chi(t)|\hat{\mu}(t)|\psi(t)\rangle), \quad (3.16)$$

$$\Delta S(t) = \frac{\Pi(t)}{\lambda} \nu(t), \quad (3.17)$$

⁵ Squaring $\Delta S(t)$ ensures that the cost never decreases when changing the pulse and that small changes are less costly.

where we have used equation 3.9⁶ and where λ is the inverse step size for each iteration.⁷ Hence, as equation 3.16 shows, the more dissimilar the forward and optimising states are at time t , the larger the update $\Delta S(t)$.

Lastly, we set the cost function from equation 3.13 as

$$\mathbb{C}(t) = \frac{\lambda}{\Pi(t)} \Delta S(t)^2. \quad (3.18)$$

Applying Krotov's algorithm to our dissipative system

When applying Krotov, we must propagate the initial state from time 0 to time T , as we did in the previous sections with Nelder-Mead. However, for the optimising trajectory we do the time-reversed process, from T to 0. Here, we must be careful: dissipation is irreversible, yet it is essential for our V-STIRAP process to work. A successful photon generation relies on the decay of the cavity field (through the mirrors), which sends the photon out of the cavity and decouples it from the system.

We must adjust the master equation 3.1. Following appendix B, we define

$$\frac{d|\hat{\rho}(t)\rangle\rangle}{dt} = \hat{\mathcal{L}}|\hat{\rho}(t)\rangle\rangle, \quad (3.19)$$

which resembles the Schrödinger equation, except now we are in Liouville space. Practically, this means we have replaced operators with superoperators (denoted with double clarets) and we have 'folded out' the density matrix into a vector (double-braced here to denote that it is in Liouville space). The dimensions of the operators and vectors have increased from $\dim(N)$ in Hilbert space to $\dim(N \times N)$ in Liouville space. We can think of equation 3.19 as

⁶ Π does not appear in the derivative in equation 3.15 because the smoothing should not, and does not, have an effect on what the pulse update should be: only what the final pulse shape is after an iteration. The krotov package takes care of this automatically for us.

⁷ Because of Krotov's monotonic convergence, there is no need to scan λ in order to find a working step size. If we need faster convergence, we may also use an adaptive λ [81].

including the Hilbert space of both the system and the part of the environment into which it decays. Likewise, $\hat{\mathcal{L}}$ describes the dynamics resulting from both the system's Hamiltonian and from dissipation. Hence, we can use it like the Schrödinger equation. More details on this are available in appendix B.⁸

Because we now work with the master equation in Liouville space, we can also take the Hermitian conjugate of $\hat{\mathcal{L}}$ when we calculate the time-reversed evolution. By doing this, the Hermitian part of $\hat{\mathcal{L}}$ will be propagated backwards in time, but the non-Hermitian dissipative part will not be: instead of reverting the decays, the system instead decays a second time [83]. Therefore, we can properly model the dissipation when using Krotov.

3.4.2 Performing pulse shape optimisation with Krotov

Before we proceed, we note that the results of this section show that, although Krotov has been applied before to many systems similar to the one we consider, Krotov does not work well for our own system. We will explain why this is the case as we continue.

Figure 3.5 shows a pictorial representation of how Krotov updates S to optimise the STIRAP process. Each step is as follows:

1. Make a forward trajectory, $|\psi_{\text{init}}(0)\rangle\rangle \rightarrow |\psi(T)\rangle\rangle$.
2. Make an optimising trajectory: $\langle\langle\chi(T)| \rightarrow \langle\langle\chi(0)|$, the time-reversed process, where we choose the new state $\langle\langle\chi(T)|$ using equation 3.14 (such that this trajectory will minimise J_T as much as possible).
3. Update the pulse S using equation 3.17 (i.e. depending on how the forward and optimising trajectories differ), and make a forward trajectory under the updated S .
4. Make another optimising trajectory.

⁸ Further details on how Krotov can be applied to open systems are available in [81].




Figure 3.5: Pictorial representation of Krotov's algorithm. The numbered steps correspond to those in the main text. In one iteration, we sequentially update the pulse S from time 0 to T in discrete time steps of length dt ; the update at time t is represented by the vertical dashed line in step 3.

5. Repeat steps 3–4 until the two trajectories match, thus converging to the minimum J and giving S_{opt} .

We now apply this optimisation technique to photon generation. Figure 3.6 shows the initial and optimised dynamics after a single iteration, and the shape of the updated pulse. We find that, when optimising **V-STIRAP**, its performance is limited. The optimising trajectory, which we stress should not necessarily be normalised (unlike the forward trajectory), displays a large increase in the populations of the states along the cavity decay path near time T . Then, the next forward trajectory suffers due to the resulting large spike at the beginning of the updated pulse. The resulting dynamics are less optimal than the initial dynamics: the transfer efficiency to $|g, 0\rangle$ is smaller. What causes this optimisation behaviour? When working with dissipative systems, Krotov always tries to minimise dissipation, which it sees as 'bad'. This is because when deriving the pulse update formula using the variational principle, the algorithm follows the gradient of steepest descent of J that does not include decays. If decays were included in the update formula, then we would need to know what decay path the system took: in the time-reversed direction, it is impossible to know which path we came from. For general open systems where we usually aim to minimise all decays, this is fine. But,

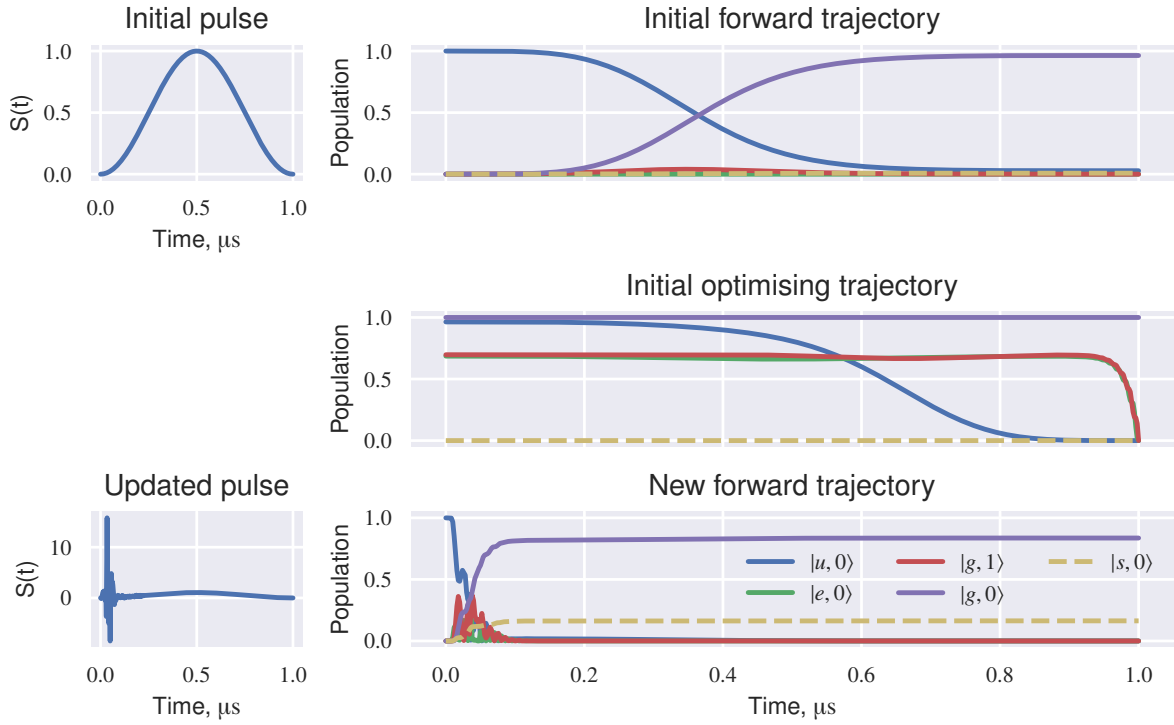


Figure 3.6: Optimising the pulse shape for photon generation using Krotov’s algorithm. We show the first iteration here only. **Top:** initial dynamics, starting with a \sin^2 pulse. **Middle:** optimising trajectory under the same pulse. **Bottom:** next forward trajectory, which uses the updated pulse shape. Here, we denote the virtual sink state to which all spontaneous emission is directed as $|s, 0\rangle$.

V-STIRAP relies on cavity decay to work. Hence in Krotov, the updated pulse fails to account for the decay appropriately, and fails to decouple the atom and cavity to complete the **STIRAP** process. Therefore, Krotov struggles with this dissipative part, and makes no changes in the non-dissipative part. Altogether, this makes Krotov, and by extension other gradient-based quantum optimal control methods such as **GRAPE**, incompatible with optimising single photon generation/absorption that relies on **V-STIRAP**, due to their reliance on the variational principle in their derivations.

Note that Giannelli et al. [93] sidestep this issue when optimising single photon absorption by setting $\gamma = \kappa = 0$ during the back-propagation and setting the target to $|g, 1\rangle$ rather than $|g, 0\rangle$.⁹ With this limitation, in the strong coupling regime the presence of a cavity photon

⁹ Kumar et al. [94] also sidestep this issue in the same way.

will cause Rabi oscillations between the states with one and zero cavity photons. So, due to the cavity coupling from $|g, 1\rangle$ to $|e, 0\rangle$, the optimal pulse by that approach will also populate $|e\rangle$, and is thus ineffective for [STIRAP](#).

Goerz and Jacobs [95] sidestep this issue in another way when optimising state transfer in a cavity network in a cascade configuration (one cavity is connected to the next, which is connected to the next, and so on). They make the assumption that, assuming the [STIRAP](#) parameters are close to optimal, the [STIRAP](#) process is always successful because the cascading coupling of cavities means the overall system remains approximately in the dark eigenstate.¹⁰ Therefore, they do not have to consider decay as part of the transfer process itself. Furthermore, because they use a quantum jump approach, they would in principle be able to switch off or on the ‘good’ cavity decay channel and ‘bad’ atomic spontaneous decay channels separately—but only because they remain in this dark eigenstate, and so could distinguish the two decay processes by whether they were accompanied by cavity emission (or absorption, in the case of the optimising trajectory) or not. While making such assumptions may work for that larger network in that particular cascade configuration, smaller, more practical networks cannot be optimised with the same approach.

With the above in mind, we will not use Krotov’s algorithm further; instead, we proceed by adjusting the Nelder-Mead algorithm to optimise the pulse shape in the next section.

3.5 Nelder-Mead pulse shape optimisation

As before, our goal is for the optimal pulse shape S_{opt} to generate/absorb photons efficiently, and without loss. Simply put, S_{opt} is the one which gives the best performance (minimises $J(T)$). But now in addition, our goal is also to maximise the coherence of the process and

¹⁰ As discussed in section 2.2.3, this means the excited state is never coupled, and so spontaneous emission can be neglected.

to control the photon's temporal length. Ensuring coherent photon generation is essential for networking, as any incoherence may result in an irreversible loss of information when we perform a remote state transfer. Likewise, the ability to change the photon length is also essential for networking, because some nodes may have a vastly different bandwidth depending on their atom-cavity parameters. For example, a node with a particular bandwidth may reach optimum performance for tailored repetition rates of single photon generation from another node.

As a result of our new goal, we now expand the definition of J from equation 3.7 to confine the photon's wavepacket to a time window and to account for the process coherence:

$$J = (1 - \eta_{\text{window}}) + \Upsilon + (1 - \beta), \quad (3.20)$$

where η_{window} is the photon generation efficiency within a particular emission window $t_{\text{start}} \leq t \leq t_{\text{end}}$ and β is the coherent component that is adapted from the standard definition of purity [96]. At time t ,

$$\mathbb{B}(t) = \left| \frac{\rho_{\text{if}}(t)}{\sqrt{\rho_{\text{ii}}(t)\rho_{\text{ff}}(t)}} \right|^2, \quad (3.21)$$

where

$$\begin{aligned} \rho_{\text{if}}(t) &= \langle u, 0 | \hat{\rho}(t) | g, 1 \rangle \\ \rho_{\text{ii}}(t) &= \langle u, 0 | \hat{\rho}(t) | u, 0 \rangle \\ \rho_{\text{ff}}(t) &= \langle g, 1 | \hat{\rho}(t) | g, 1 \rangle, \end{aligned}$$

with the subscript i or j referring to the initial or final state of the **STIRAP** process respectively. Thus, $\mathbb{B}(t)$ is a ratio between the off-diagonal and diagonal elements of $\hat{\rho}$: alternatively, the ratio between its coherences and populations. If the transfer stays fully coherent, $\mathbb{B}(t) = 1$;

if fully incoherent, $\mathbb{B}(t) = 0$. To calculate how coherent the photon emission is over the emission window, we use

$$\beta = \frac{1}{\tau} \int_{t_{\text{start}}}^{t_{\text{end}}} \mathbb{B}(t) dt, \quad (3.22)$$

where $\tau = t_{\text{end}} - t_{\text{start}}$. As before, we use a sink state in our simulation, ensuring that we only maximise the coherent transfer with our optimisation. To illustrate the coherence, in figure 3.7a we plot \mathbb{B} for the simulation we showed in figure 3.1b where we used a \sin^2 shape and the optimal atom-cavity parameters.

After our new definition of J , we must construct a representation and associated parameters for S itself. What should these be? Our requirements are that the pulse is:

- smooth, meaning that $S = dS/dt = 0$ at the pulse extremities;
- adiabatic/varied appropriately slowly, meaning that dS/dt is kept well below the energy splitting that results from the atom-cavity interaction [26];
- constrains the photon emission to a window of length $\tau = t_{\text{end}} - t_{\text{start}}$.

Taking all of this into consideration, we choose to represent S by a sequence of equally-spaced amplitude points with a spline fit through them, as shown in figure 3.7b. The amplitudes are our parameters, and the spline fitting ensures that the pulse is smooth. We use $N = 21$ points, striking a balance between computation time and a finely-enough controlled time window. Additionally, to ensure that the pulse starts and ends smoothly, we keep the pulse endpoints fixed at zero (so we only optimise 19 pulse points) and we again use a smoothing/window function $\Pi(t)$ of half- \sin^2 shape, such that

$$\Omega(t) = \Omega_0 S(t) \Pi(t). \quad (3.23)$$

We choose from two different kinds of splines: a smoothed spline, and an interpolating spline.




Figure 3.7: How we use coherence and splines with Nelder-Mead to optimise pulse shapes. **(a)**: coherence for optimal atom-cavity parameters with a \sin^2 pulse. Following equation 3.22, we integrate \mathbb{B} with respect to time to get the coherence β : that is, we calculate the area under the curve. Here, $\beta = 99.4\%$. **(b)**: Initial \sin^2 shape for the Nelder-Mead pulse shape optimisation, and how we fit splines through it. Blue points denote the amplitudes that we will optimise. We fit a smoothed spline (red dashed) through the points, unless we are optimising for emission in a particular window in which case we fit an interpolating spline (green dot-dashed) that goes through every point.

Smoothed splines may yield a more efficient **V-STIRAP** process with less computational effort, as the ‘knots’ of the spline (the points which the spline passes through) have some leeway to move to earlier or later times. However, for a shorter emission window/photon, we may want to fit through a few of the central pulse points and omit the rest, but then we would have to weight the knots appropriately. To sidestep having to weight the knots, we instead use an interpolating spline when optimising emission in a given time window, and a smoothed spline for all other optimisations.¹¹ With Π and with $N = 21$, we can control τ from 1 to $0.1 \mu\text{s}$.

As an aside, although using a Fourier decomposition to represent S instead of interpolated amplitudes seems possible, it is much less efficient here because it requires a much larger N [61], especially for τ control.

¹¹ From the `scipy` Python package, we use `UnivariateSpline` with smoothing parameter $s = 1$ for the smoothed spline, and `CubicSpline` with the ‘clamped’ boundary condition parameter for the interpolating spline (meaning the two endpoints are not knots, and remain fixed to zero amplitude). We set both to use polynomials of degree 3 (i.e. cubic) to avoid oscillations resulting from numerical instability [97].

We firstly apply the Nelder-Mead algorithm to the pulse amplitudes to optimise the pulse shape, not constraining to a particular time window. Figure 3.8 shows S_{opt} and the resulting dynamics for optimised photon generation and absorption. η_{gen} has now further improved from 95.4 to 99.1 %, η_{abs} from 39.4 to 62.8 %, and \mathcal{V} from 6.6 to 0.8 % in generation and reaching 1.8 % in absorption. so the process is both efficient and with a high fidelity. We see that the pulse shapes are close to what we may expect from the literature: asymmetric, and gradually ramping up near the end(beginning) of the generation(absorption) [59, 75, 76, 98].

In addition, the absorption coherence of the **STIRAP** process reaches $\beta_{\text{abs}} = 99.6$ %, also highlighting the high process fidelity. The generation coherence reaches $\beta_{\text{gen}} = 98.1$ %. Although high, β_{gen} has decreased slightly from the initial pulse (down from 99.4 %). If we look at the emitted photon shape, we can see why: most of the emission is at the beginning of the process, and the photon has a long tail with negligible emission/cavity population. This small fraction represents the ‘finite support’ that the **STIRAP** pulse offers [59]: because transferring population from the atom into the cavity gets more and more difficult as the atomic population gets smaller, then we would need an infinitely long pulse to transfer 100 % of the population. So, because in the laboratory we must use pulses of finite length, we therefore restrict the simulated process to take place between time 0 and T —a finite support—leaving a remaining negligible population. If achieving coherence was our priority over efficiency, we may therefore add a small initial population in $|e, 0\rangle$ to account for this and achieve better coherence. What we find particularly notable is that this matters for generation as well as absorption: in [59], only absorption was considered.

Returning to the efficiencies, for generation we note an overshoot of $\Omega(t)$ above Ω_0 , to about $1.2\Omega_0$. Although a stronger peak driving may be expected to be a contributor to the η_{gen} increase, we instead draw our attention to the gradient dS/dt in the central region of the pulse: S_{opt} has a more gentle gradient than that of the initial \sin^2 pulse. Thus, this suggests

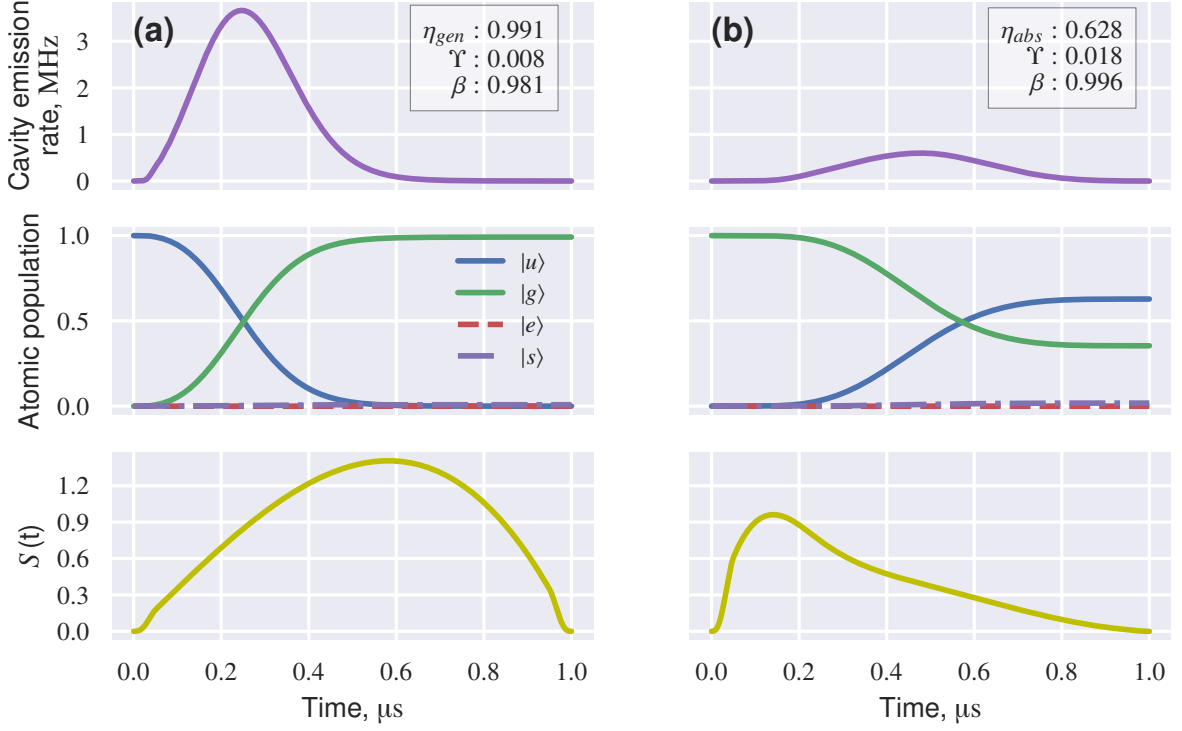


Figure 3.8: Optimal laser pulse shapes for single photon emission and absorption. We use the optimal atom-cavity parameters (see table 3.1), and set Raman resonance for the absorption in particular. Referring to equation 3.23, we multiply the optimal pulse shapes (yellow) by Ω_0 to obtain these dynamics.

that S_{opt} better meets the adiabatic condition in this regime of strong driving, and is likely to be the main contributor to the η_{gen} increase.

For absorption, we note that the particularly large improvement in η_{abs} may be due to the resulting absorption pulse giving a reduced cavity emission. Its maximum $S(t)$ is lower than that of the generation pulse. This suggests that impedance matching is particularly important when considering photon absorption efficiency, and may also explain why the generation reaches a higher efficiency. As such, this result highlights the importance of driving the system with the optimum pulse power Ω_0 in order to have an impedance-matched in-coupling to the cavity.

Next, we investigate how we can control the length of emitted photons using the optimisation.

Figure 3.9 demonstrates how as we gradually narrow the emission window τ , the algorithm compensates by narrowing the pulse. Since shorter photons require stronger/faster driving pulses, we now bind S to the range $[0, 2.0]$.¹² We run this optimisation for 12 values of τ from 1.0 to 0.05 μs : figure 3.9a shows the corresponding pulse shapes for two of these values, and figure 3.9b shows the efficiencies for all 12 τ values. Interestingly, we see in figure 3.9a that for shorter pulses, the optimisation converges to a multi-pulse V-STIRAP process. Most of the population transfer takes place in the first, fast pulse section at t_{start} , after which the following pulse section(s) transfers the remaining population. This keeps the photon emission highly efficient, with only a small component outside of the emission window, until the required driving is too fast/non-adiabatic. Our intuition on why multi-pulse V-STIRAP may be more efficient focuses on the vanishingly small cavity population that remains after a single V-STIRAP pulse. The less cavity population remaining, the harder the system must be driven to retrieve it, due to the fixed g and κ . Using multiple weaker pulses rather than a single strong pulse may therefore be more efficient overall as we intrude into the non-adiabatic regime, where the efficiency punishment for using too strong a laser pulse is more severe.

Seeing from figure 3.9b the point where η_{window} starts to fall and the process becomes non-adiabatic, we therefore identify the region with a quantifiable tradeoff between η and τ . This suggests that the bandwidth of our photons is limited by the efficiencies we demand. Thus, in cavity-based quantum networks where nodes have different bandwidths due to their differing spectral resolutions (i.e. cavity linewidths κ), one must take note of this tradeoff between bandwidth and efficiency. With our optimisation, we can account for this tradeoff.

Note that the optimal parameters for the system considered here (see table 3.1) are well into the strong-coupling regime, but that marginal losses still occur on a similar timescale to

¹² Note that we bind the pulse points and not the interpolating spline itself: this is why figure 3.9a shows $S(t)$ reaching slightly above 2.0, as the spline must fit through two successive points that are already close to 2.0.

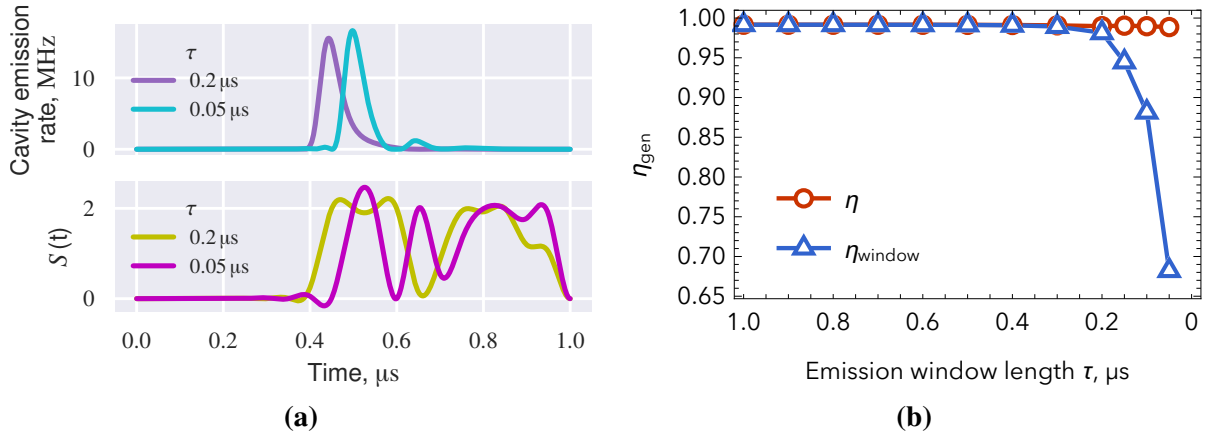


Figure 3.9: Optimising S for emitting photons efficiently within a particular time window. **(a)** How S_{opt} and the emitted photon shape varies as we reduce the desired photon length τ . We plot two shapes corresponding to $\tau = 0.2 \mu\text{s}$ and $\tau = 0.05 \mu\text{s}$. Note that we are using interpolating splines rather than smoothed splines here. **(b)** Dependence of emission efficiency on the desired photon length τ . The overall efficiency η does not fall greatly, but η_{window} starts to fall as the STIRAP process becomes non-adiabatic.

the length of the driving pulse. This matches the explanation of Giannelli et al. and Macha [61, 93] where numerical optimisation techniques are found to perform more effectively when close to the adiabatic limit and in the non-adiabatic regime than the analytic solutions. We postulate that this is due to the increasing difficulty of maintaining coherence as the process itself becomes longer, as our coherence results suggest. Furthermore, this supports our assertion that our optimisation approach compensates well for the loss at the transition between the two regimes.

One possible modification of our Nelder-Mead approach is to change the optimisation's basis. To be precise, instead of optimising pulse amplitudes on a coarse time grid, we could optimise the characteristics related to one specific photon shape, e.g. Gaussian, that is well-suited for networking.¹³ We would then parametrise the asymmetry, length, area, leading and trailing edge heights etc. of the driving pulse, and optimise those. The shape of

¹³ Gaussians are well-suited network shapes because they are symmetric, smooth, and easier than other shapes to create in the laboratory.

the resulting photon would then vary accordingly. In this way, the photon shape could be optimised for a common shape for all nodes. In the literature, this basis change is analogous to using the chopped random basis (CRAB) algorithm [99, 100]. However, this basis definition then forgoes some control over the pulse shape. Since our results suggest that pulse shape control is paramount to minimising the remaining losses in this process, it is speculative whether this approach would improve upon the current one.

To summarise, after optimising the **STIRAP** laser pulse shape using the Nelder-Mead algorithm, we achieve an efficiency and coherency of $\eta_{\text{gen}} = 99.1\%$ and $\beta = 98.1\%$ respectively for photon generation and $\eta_{\text{gen}} = 62.8\%$ and $\beta = 99.6\%$ for photon absorption.

3.6 Optimising remote state transfer

Having optimised single photon generation and absorption separately, we now bring them together. We connect two cavities to form an elementary quantum network, as previously shown in figure 1.1b. An optical fibre forms the direct connection. To alleviate the effect of reflections within this fibre interfering with the photon wave packet being sent, we insert an optical circulator (not shown in figure 1.1b): a device with three inputs labelled A, B and C which directs light from A to B only, B to C only, and C to A only. This way, light from one cavity (A) may only travel towards the other (B).

We emit a photon from one cavity and absorb it in the other, following the remote state transfer scheme in [23] that we mentioned in section 3.1 in which we map between the atom's electronic state and the photon's polarisation state. Each cavity has identical parameters (the optimised ones in table 3.1). Since here we are not concerned with polarisation-dependent effects that have already been well-reported elsewhere [32, 101, 102], we only use the initial and final $|F, m_F\rangle$ states defined earlier in section 3.3,¹⁴ and so our photons are exclusively

¹⁴ Recall that we only simulate 4 atomic levels here.

left-handed circularly polarised (exciting $\Delta M_F = 1$). In our simulations, sending the photon is straightforward: we take the emitted photon shape from the generating cavity and feed it as a photon in-coupling in the absorbing cavity. By definition, this wave packet is already normalised, containing an average photon number η_{gen} . We also assume that the fibre is fairly short (~ 10 m) and so we do not consider the effects of fibre loss here. The consideration of fibre loss on similar cavity-based state transfers and networking schemes can be found for example in [24, 103, 104].

With the system set up, we must modify J for this optimisation. We simply combine our performance functions for generation and absorption to define

$$J = (1 - \eta_{\text{gen}}) + (1 - \eta_{\text{abs}}) + \Upsilon_{\text{total}} + (1 - \beta_{\text{gen}}) + (1 - \beta_{\text{abs}}). \quad (3.24)$$

Note that our choice to calculate the sum of these terms rather than the product allows us to more simply weight individual terms if we so choose, to determine how important each term is to the performance; although, we did not weight any terms, in order to keep our optimisation simple. We start with \sin^2 shapes as before and then optimise $N = 2 \times 19 = 38$ pulse points using Nelder-Mead.

Figure 3.10 shows the optimal pulse shapes and dynamics for this elementary network with identical cavities, yielding an overall state transfer efficiency $\eta_{\text{gen}} \times \eta_{\text{abs}} = 39\%$. The optimal shape for the generation is similar to before with its slight asymmetry, and gives the same dynamics. On the other hand, the absorbing pulse is less efficient than for the standalone absorption optimisation, down from 62.8 to 39.4%. This drop can be explained by looking at the emitted photon shape from the generating cavity: as the amplitude of the emitted wave packet gets larger, the absorbing pulse ramps up to compensate by keeping the cavity emission negligible, as we would expect. However, once the amplitude becomes too large, the absorbing node fails to do so, and so the absorbing node's emission rate starts to increase. The

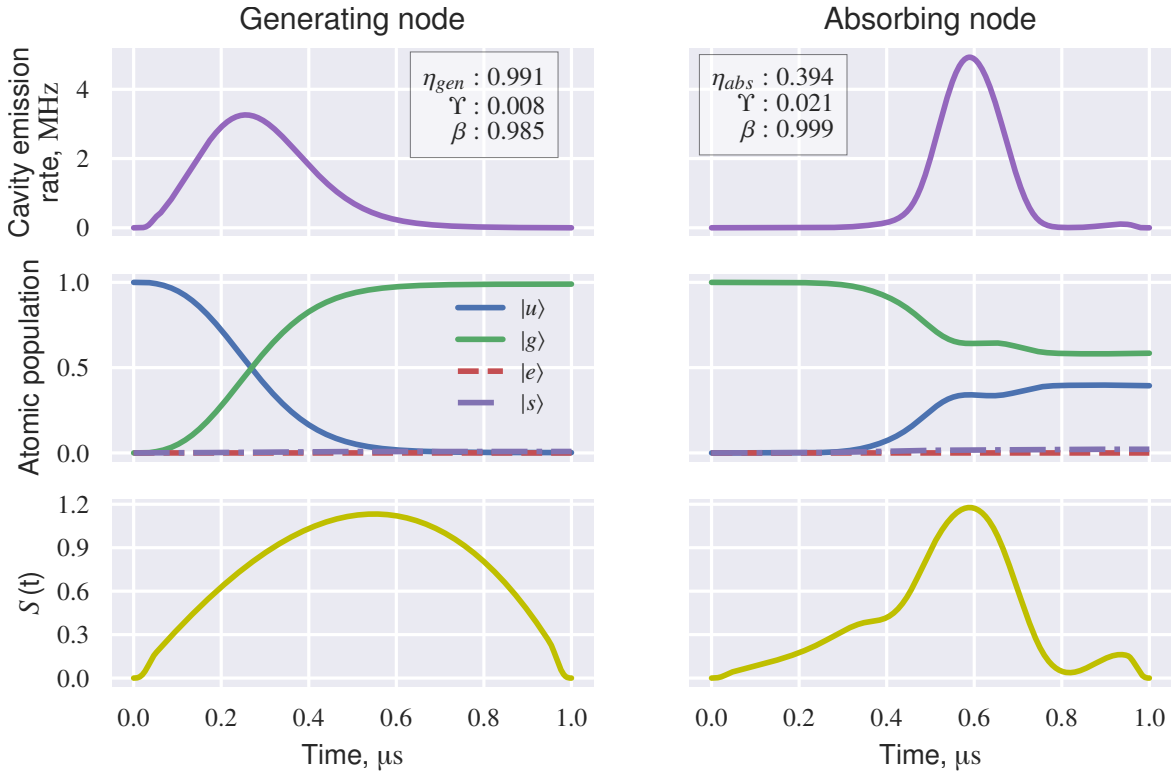


Figure 3.10: Optimal laser pulse shapes and resulting dynamics for a remote state transfer between two cavities. Both cavities have identical parameters. We send the emitted photon from the generating cavity into the absorbing cavity. Note that we set the absorbing node on Raman resonance, as before.

time at which this happens coincides with the kink in the absorbing node’s atomic population transfer. Therefore, the presence of this kink suggests how important it is to find the optimal linewidth κ for transfer between cavities: the optimal κ for emission and absorption in a single node may not be the optimal κ for cavity-based networking. Following our discussion in section 3.5, we could think of a large κ as allowing the photon to escape from the emitting cavity easily, but requiring a stronger laser pulse in the absorbing cavity—which could incur an efficiency penalty since the (optimal) process may be close to the non-adiabatic regime.

Note that our optimal pulses do not match the model by Cirac et al., wherein for this case of identical cavities, the absorption pulse would have the same shape as the emission pulse shape but reversed in time [25]. This may be because in that work, they derive pulses for

emitting and absorbing time-symmetric photons, while here, we find that the most efficient generations and absorptions are not time-symmetric.

Motivated by our previous result, we further investigate the effect of κ on our state transfer by considering the directionality of our cavities. As we mentioned in section 3.2, a realistic cavity has two input and outputs: the two mirrors. We have been modelling cavity emission along only one direction. In many experiments, the cavities are directional (one mirror is more reflective than the other) to allow efficient emission of a cavity photon [21]. Yet, some cavities may be symmetric (both mirrors have the same reflectivity), such as the fibre cavity we will assemble in the later chapters. To model this scenario, in figure 3.11 repeats the state transfer optimisation, but with directionality modelled (using $\kappa_{\text{in}} = 0.35\kappa$, as mentioned in section 3.2). Now, the transfer efficiency is lower, at $\eta_{\text{gen}} \times \eta_{\text{abs}} = 20\%$. Compared to figure 3.10, we see that the cavity emission from the absorbing node has dropped by a factor of about 5, yet η_{abs} has only dropped by a factor of 2. This suggests two things. Firstly, the difficulty of coupling light into the absorbing node is the reason for the transfer efficiency drop. Secondly, because we have a smaller incoupling rate ($\kappa_{\text{in}} < \kappa$), we can more easily impedance match the incoming photon to keep the efficiency high for a typical experimental setup. Thus, our result highlights the importance of choosing the right pairing of mirrors when performing remote state transfer. In future simulations, we can also extend our model to include the input/decay channels from the left and right cavity mirrors separately, allowing us to investigate an ideal symmetric cavity with our master equation approach [105].

The flexibility of our numerical approach allows us to optimise state transfer for non-identical cavities too: a more realistic case, since every cavity is slightly different due to manufacturing tolerances. Here, we optimise a hybrid cavity network in figure 3.12. To match our laboratory scenario, we take the generating node to have the optimal parameters from before (our fibre cavity), and the absorbing node to have the parameters of the macroscopic

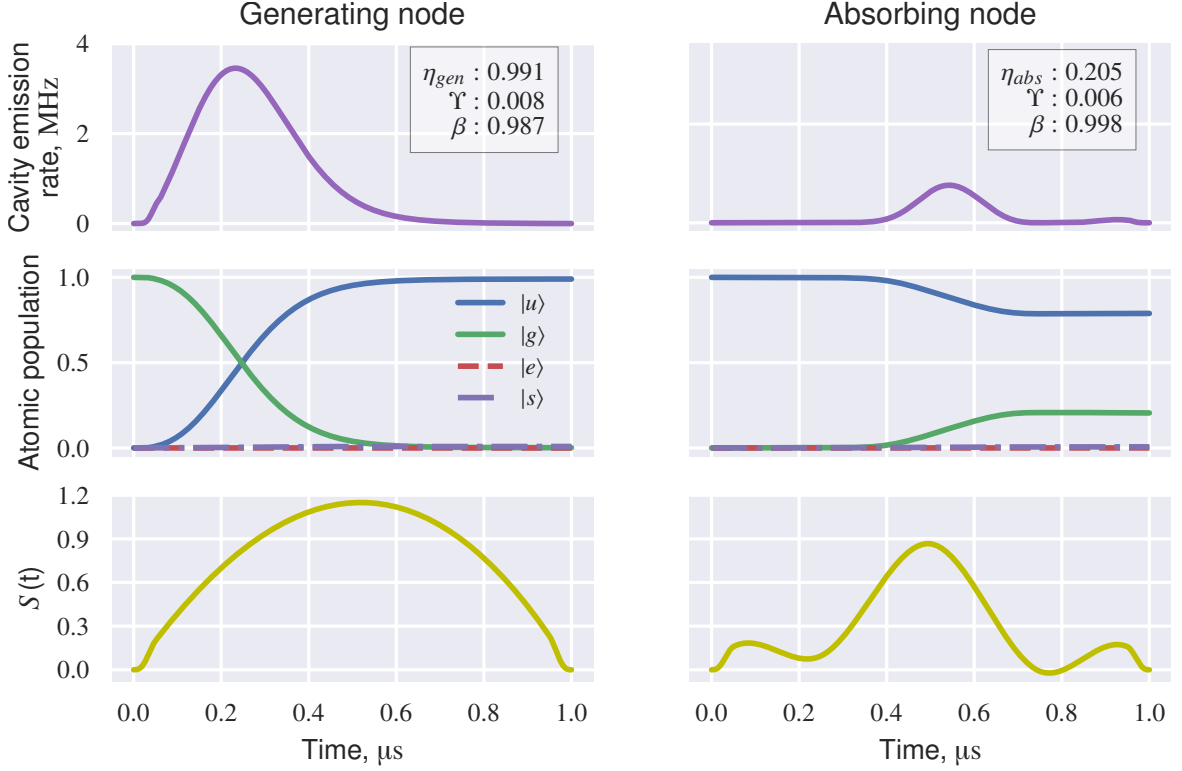


Figure 3.11: Optimal laser pulse shapes for remote state transfer, with transmission and loss channels modelled using equation 3.3. Compared to figure 3.11, we now model the absorbing node with input coupling rate $\kappa_{\text{in}} = 0.35\kappa$.

cavity used by Barrett [32, 55].¹⁵ These parameters are $(g, \kappa, \Omega_0)/2\pi = (4.8, 3.8, 8.7)$ MHz and $\Delta_L/2\pi = \Delta_C/2\pi = 52$ MHz. We assume $\kappa_{\text{in}} \approx \kappa$ for the absorbing cavity, which is directional. As we would expect, especially since Barrett’s directional cavity was designed for single photon emission and not absorption, the transfer for this hybrid setup is less efficient than before, at 11 %. Most of the light the absorbing node appears to couple into the cavity but then leak back out. Moreover, it emits a even larger amount of light than the generating node. Judging from the low $\beta = 69$ %, the incoming light may be coupling to the absorbing atom before being re-emitted. This is also supported by the fact that S at the generating node has a lower maximum amplitude than before, suggesting that for such a network to work

¹⁵ The parameters are available in [55, table 3.1].

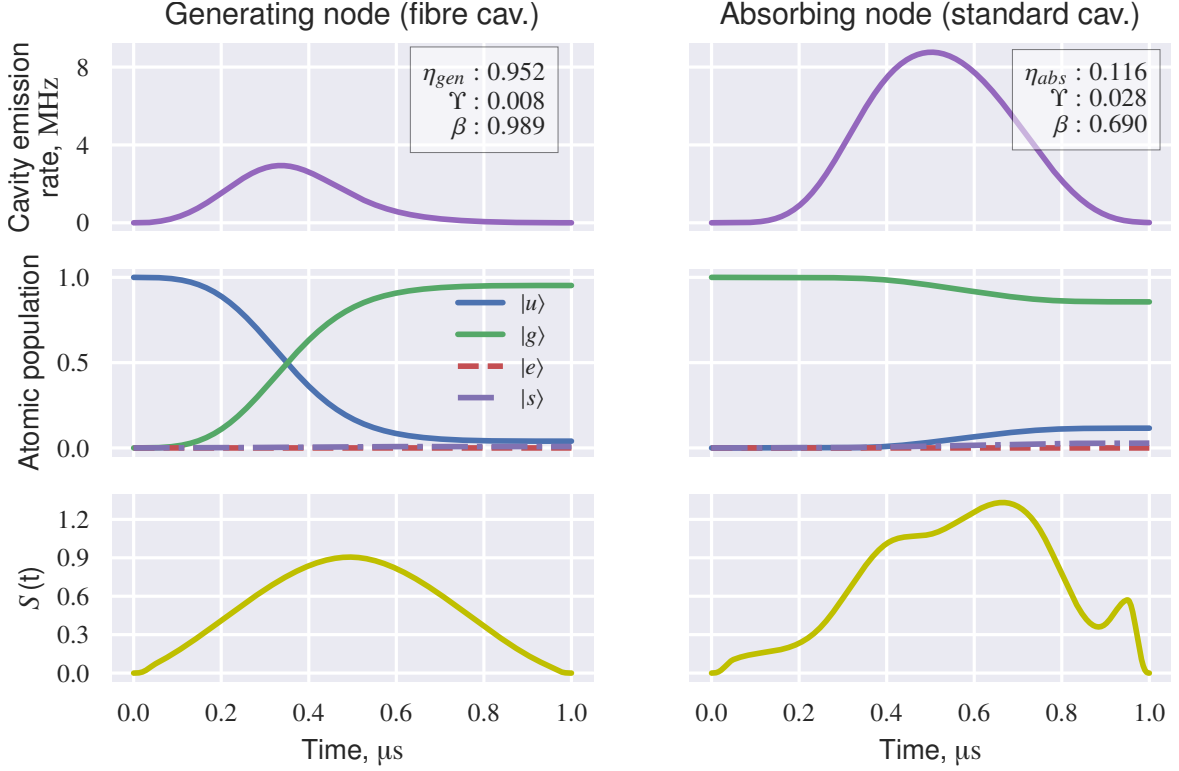


Figure 3.12: Optimal laser pulse shapes for remote state transfer in a hybrid cavity network. Each node has different parameters: the generating node has our optimal parameters found earlier, and the absorbing node has the parameters from [55], which are $(g, \kappa, \Omega_0)/2\pi = (4.8, 3.8, 8.7)$ MHz and $\Delta_L/2\pi = \Delta_C/2\pi = 52$ MHz. We stress that the Ω_0 used for the two cavities are different.

well, one must compromise the generation process in order to maintain reasonable absorption efficiency with smaller g , thereby maximising the possible transfer efficiency.

3.7 Conclusion

We have applied numerical methods to optimise the parameters for a quantum network of atoms in cavities. Specifically, we have applied the Nelder-Mead algorithm to optimise both the node parameters and the [STIRAP](#)-driving laser pulse shape used to generate/absorb a single photon from/into a node. For parameters corresponding to a fibre-tip Fabry-Pérot cavity in the strong coupling regime, we obtain a photon generation efficiency η_{gen} of 99.1 %

and a coherence β of 98.1 %, and absorption efficiency η_{abs} of 62.8 % with 99.6 % coherence. This is in line with or above that achieved when applying other optimisation schemes [59, 75, 76], and will prove useful for extending the performance of network operations such as remote entanglement [23, 65] or storage in quantum memories [16].

We have also applied the Krotov algorithm to optimise the pulse shape for the same system. Due to the reliance of the considered **V-STIRAP** process on irreversible cavity emission, we must modify the algorithm to suit the photon generation and absorption processes. We find that Krotov's algorithm, and by extension other gradient-based algorithms which minimise dissipation, are therefore poor choices for optimising processes that use **V-STIRAP**, such as the cavity-based quantum network of Ritter et al. [23].

We have optimised the remote state transfer of a photonic qubit between two cavities, wherein the photon generated by one cavity is absorbed in the other. For two identical fibre cavities, we achieve a transfer efficiency of 39 %. When introducing the directionality of emission into our model, we achieve 20 % efficiency, highlighting the importance of choosing suitable mirror reflectivities for the left and right mirrors. From this, we have found in general that what is optimal for generation or absorption alone is not optimal for state transfer. In the same vein, we have found that the most efficient generation and absorption pulses in all cases are not necessarily time symmetric. We further optimised a hybrid cavity network containing one fibre cavity and one macroscopic cavity, achieving a lower efficiency of 11 % due to the input light coupling in to the cavity and then being re-emitted by the atom before escaping again. By achieving this, we have highlighted that one of the main impediments to performing efficient remote state transfer in hybrid networks is coupling the emitted photon to the absorbing node, which improves with better impedance matching of the input photon and cavity fields. Likewise, we found that keeping the incoupling rate below the cavity linewidth gives more leeway in improving impedance matching.

Given our simulations' modelling of a specific level structure for use as our network's stationary qubit, our approach can be extended to other network implementations. One would replace the atom with the relevant stationary qubit, such as ions or colour centres, by adjusting the simulated level structure appropriately. For this reason, we heuristically expect that our findings relate to other network implementations in general. Altogether, we have gained essential insights towards, to quote Kimble, 'whether a quantum network "works" ' [14].

Making a Fibre-tip Fabry-Pérot cavity for fully deterministic quantum interfacing

From the previous chapters, it is clear that the cavity is the heart of our research. In chapter 1, we mentioned how to realise a quantum interface between cavity photons and atoms trapped inside. In chapter 2, we explained how this cavity makes the atom-photon interaction stronger, making interfacing possible. In chapter 3, we simulated this to see what capabilities an optimal interface would have. Therefore, the question arises: how do we make this interface in the laboratory?

Making our interface in the way we envisage requires a small mode volume for strong coupling, and enough optical access for a tightly focused beam with a micrometre-sized waist. Thus, we must use a microcavity, as a short cavity with larger mirrors will block the beam. We also want a pre-aligned cavity in vacuum, rather than an actively aligned cavity requiring a bulky translation stage that could block the beam and also reduce cavity stability. In addition, efficient coupling of light in/out of the cavity is paramount. We are therefore motivated to make a [fibre Fabry-Pérot cavity \(FFPC\)](#)¹ to satisfy these constraints. Figure 4.1 shows our envisaged interface. Not only can we trap one atom at the cavity waist, we can also combine the tweezer light with a spatial light modulator (SLM) to trap a 2D array of atoms

¹ We will use the terms [FFPC](#), fibre cavity, and cavity interchangeably from here.




Figure 4.1: Our envisaged fibre cavity atom-photon quantum interface. With enough optical access (large numerical aperture (NA)), we can move single atoms into the cavity mode using optical tweezers without the tightly focused tweezer beams clipping the fibres. Each tweezer individually manipulates a single atom.

and move them independently of one another [106–108], as we will discuss in chapter 5 and appendix C.

Our challenge is to assemble the FFPC precisely enough to meet these constraints whilst optimising the interplay between maximal atom-cavity coupling strength, efficient in/out-coupling and optimal optical access. To this end, in this chapter we develop and demonstrate an assembly process for a fibre cavity-based atom-photon interface that is permanently aligned with sub-micrometre precision—no active alignment with a translation stage is needed like in other implementations [66, 109, 110]. We note that the methods devised here apply generally to similar devices that require precise positioning: for example, colour centre interfaces in a cryostat [111, 112] or ion interfaces in vacuum [113].

In overcoming this challenge, we make the first FFPC formed of two single-mode (SM) fibre-tips in vacuum. All other implementations use at least one multi-mode (MM) fibre-tip in order to alleviate the demanding alignment requirement. Thus, as both mirrors we will use are equally reflective, our successful assembly shows that it may indeed be possible to realise fully symmetric operation for the quantum networking scheme envisaged by Kyoseva et al [114]. It is this symmetric operation that evokes the phrase ‘fully deterministic’ in the

title of this chapter since any, even small, amount of light escaping through an MM fibre would fundamentally count as lost because of mode-mixing in the fibre [115, 116]. Our device eliminates this fundamental loss, as we will discuss in chapter 6.

In achieving this goal, we firstly describe how fibre-tip mirrors are made in section 4.1, so that we can understand how our cavities will perform. Next, we establish the mode-matching theory that quantifies what alignment precision we need in section 4.2, as well as the effects of mirror manufacturing tolerances on our cavity design choices in section 4.1. Then, we give our FFPC assembly methods in section 4.4, followed by the results and discussion of our final assembled cavity in section 4.5. Finally, we conclude in section 4.6.

4.1 How are fibre-tip mirrors made?

We already know from section 2.1.2 that a stable FPC requires curved mirrors. Practically, an FFPC simply consists of two optical fibres aligned towards each other, with the curved mirror surfaces implemented directly on the fibre-tips. Figure 4.2 summarises how to construct it: effectively, the user cuts a fibre in half, ablates² the ends to form the curved feature, mirror-coats them, and finally aligns them.

There are two main benefits of using laser ablation to make the curved mirrors.

Firstly, using laser ablation allows us to access a much smaller radius of curvature, ζ , than with typical high-finesse cavities that use superpolished glass substrates.³ Thus, the mode volume V_m is smaller and the atom-cavity coupling g is higher.

Secondly, since here the cavity mirrors are the fibre-tips themselves, the optical access is much better than those superpolished mirrors because the fibre diameter is comparatively small, at 125 μm versus ~ 5 mm. Optical access is precisely what motivated Reichel et al. to

² Here, ‘ablation’ means removing material from the fibre-tip by vaporising that material with a laser pulse.

³ Superpolishing is simply polishing where the final surface is particularly smooth: typically surface roughness of sub-nanometre order.




Figure 4.2: Microscope images illustrating the stages of construction of an FFPC. (a) An optical fibre is cleaved to form a flat surface. (b) This fibre-tip is ablated with a high-powered CO₂ laser pulse to form a Gaussian depression. (c) The fibre-tip is mirror coated: a dielectric multilayer stack is deposited onto the fibre-tip to form the reflective surface. (d) Two such fibre-tips are aligned to form an FFPC. (a) and (b) reproduced from [125].

create the first FFPC in 2006 [117], and also the high-finesse FFPCs upon which all current implementations are based in their seminal paper [118]. It is quantified in terms of numerical aperture,

$$NA = \sin \theta,$$

where θ is the maximum half-angle that focused light can make with the optical axis (that is perpendicular to the cavity axis), as shown previously in Figure 4.1. Our FFPC affords optical access up to $NA = 0.6$: three times stronger than with superpolished mirrors, and more than enough for our optical tweezers.⁴ Additionally, we can couple cavity light directly to the fibre. However, we must accept a compromise in the maximum cavity-to-fibre coupling efficiency; we will address this in the next section.

As an aside, we note that although research has been constantly shrinking V_m for many years [21, 119–122], the mirror reflectivity R in the optical domain has plateaued since 1992 [55, 123] as the methods of producing high reflectivity, low loss mirrors has not changed [124, Ch. 7]. Therefore, shrinking V_m should be the primary way to increase g and therefore the interfacing efficiency, as the way of fabricating FFPCs presented here does.

⁴ We shall discuss the precise requirements for the tweezers in chapter 5 and appendix C.

Parameter		Value	Standard deviation	Remarks
Fibre-tip geometry				
Fibre diameter		125 μm		As specified
Fibre mode waist	w_f	3.6 μm		
Radius of curvature	ζ	220 μm	50 μm	From WLI [125, 126]
Feature diameter	D_ζ	40 μm	4 μm	
Feature depth		2 μm	0.6 μm	
Ellipticity		0.9	0.1	
Mirror coating				
Transmission	T	10 ppm		As specified by ATFilms [125]
Scattering loss	\mathcal{L}_S	10 ppm		
Absorption loss	\mathcal{L}_A	2 ppm		

Table 4.1: Specification of the fibre-tip mirrors that are used to build the FFPCs. The measured data and spreads are for 72 fibre-tips; of these, our group kept 48 and the Paris group kept 24.

For the FFPCs we use, the fibre-tip mirrors used were originally manufactured in a collaboration between former group member Jian Dong and Jakob Reichel’s research group [125, 126], with the mirror coating performed by ATFilms. Table 4.1 summarises their specification and measured data from that collaboration.⁵ The most notable parameter is ζ : demonstrating what the capabilities of these FFPCs are, for $\zeta = 220 \mu\text{m}$, $L_{\text{cav}} = 100 \mu\text{m}$ and $\mathcal{F} = 140\,000$, and using equations 2.12, 2.16, 2.54 and 2.55, we calculate a maximum possible $C = 180$ that is well into the strong coupling regime.

Ablation process

Key to the performance of the mirror is the ablation process, since this determines how smooth—and ultimately, reflective—the mirror coating can be. The ablation laser is a high-powered CO_2 laser of wavelength $11.6 \mu\text{m}$ that is focused on the fibre-tip. Since the fused silica fibre strongly absorbs light at this wavelength, the laser vaporises some material

⁵ In addition to the SM fibre-tips, a small number of flat (non-ablated) mirrors and MM fibre-tips were made, neither of which were available for the research in this thesis.

on the fibre-tip. The laser pulse length and power determines the diameter D_ζ and the feature depth; the user must take care not to melt the fibre-tip with too powerful a pulse. The laser, having a Gaussian intensity profile, leaves a corresponding Gaussian depression on the fibre-tip. Close to the centre of the ablation, the Gaussian shape is also approximately spherical [118]: this gives the effective usable diameter D_ζ . The resulting mirror smoothness is comparable to that of superpolished mirrors (surface RMS roughness 0.1 nm), so mirror losses due to scattering are low too. With respect to this, the beam's profile should not deviate from a 'perfect' Gaussian.⁶

Mirror coating

The mirror coating itself is a dielectric Bragg stack: a stack of alternating high and low refractive index thin film coatings. Each film must have $\lambda/4$ thickness so that successive reflected wavefronts constructively interfere: this explains the stringent smoothness requirement for the mirror surface. The dielectric stack is coated on the fibre-tip using ion beam sputtering (IBS). Somewhat similarly to adding metallic paint to a car, IBS coats the fibre-tip with an ionised gas, and then sputters an ionised beam of the mirror coating material(s) over the fibre-tip, leaving a thin film(s) of uniform thickness. The mirror reflectivity R should ideally be as high as possible before scattering and absorption losses start to dominate over reflections. Since the mirror roughness is proportional to the mirror losses, it defines a practical limit on R .

Furthermore, table 4.1 shows that the ablation manufacturing process naturally exhibits some deviation from the desired mirror feature. As a whole, these manufacturing errors can be broken down into three aspects: decentration, ellipticity, and clipping loss. These will be discussed later in section 4.3.

⁶ Alternatively, a mirror shape that is a solution to the Helmholtz equation may be better for cavities with ultra-small $\zeta \sim 1 \mu\text{m}$ [122, 127, 128]. This requires shaping by focused ion beam (FIB) milling rather than laser ablation, and so tends not to be considered due to the difficulty in milling fibre-tips. For us, the resulting cavity would be too short for decent optical access.




Figure 4.3: Mode mismatch between (1) single-mode fibre core mode, with radius w_f ; and (2) cavity mode, with radius w_m at the mirror/fibre-tip. On the right, we illustrate the intensity of the beam as we look at it head-on.

Therefore in the next section 4.2, we theoretically quantify what alignment tolerance we need for a decent mode-matching with our fibres. Afterwards in section 4.3, we can determine how the different manufacturing parameters of the fibre-tip mirror fabrication process affect our tolerance, and then finalise our design for the cavity.

4.2 Mode-matching

When considering the coupling efficiency of cavity light to the fibres, we want the spatial profiles of the fibre modes and cavity mode to overlap well. This is the mode-matching condition, as illustrated in figure 4.3. Both modes have Gaussian transverse profiles. Any non-overlapping fraction of these profiles is lost, since the light has nowhere to go.

Looking closer at figure 4.3 highlights our resulting mode-matching problem: the fibre mode waist w_f is smaller than the cavity mode waist at the mirror w_m , and so a significant fraction of light is lost. What fraction depends on the fibre core size itself, but also on the cavity alignment: the worse the alignment, the worse the mode-matching, and if the fibre core is smaller then the cavity is even more sensitive to misalignment. Although losing light




Figure 4.4: Modes of a misaligned cavity. Here, we misalign the cavity transversely by d and angularly by θ . We take the overlap integrals at the base of the ablated feature.

does not affect the attainable fidelity of quantum interfacing with our device [21], it certainly affects how efficient it is. The worse the mode-matching, the lower the rate of successful interfacing operations, and if too inefficient then we lose the ability to control the interfacing before the system succumbs to losses. The mode-matching issue is clearly critical to the interfacing.

To tackle the mode-matching problem, we will develop in section 4.4 a cavity alignment method that incurs the lowest reasonable loss in mode-matching efficiency. Our method, in essence, breaks down the alignment process into as small steps as possible while still maintaining full control over the final positions of the fibres. We do not need to apply any alignment feedback, thus removing those systematic errors as far as possible too.

In order to benchmark our alignment performance, we need to quantify what mode-matching we expect to obtain and, most importantly, how much further misalignment we can tolerate. Here, we use overlap integrals [34, 129, 130] to calculate both the mode-matching efficiency, ϵ , and how misalignment affects it.

We calculate ϵ following the analysis by Gallego et al [129]. The methods described herein are specific to FFPCs, for reasons which will soon become clear. For brevity, we only

reproduce the salient parts of the analysis; full details are available in [129, 131].

Firstly, we separate the mode field into an amplitude component and spatial component, exactly as we did in equation 2.24 but with a different notation:

$$E = E_0|E_r\rangle, \quad (4.1)$$

where the Dirac notation for the spatial part denotes the integral of this classical field⁷ over the 2D xy -plane at position z ,

$$|E_r\rangle = \iint_{-\infty}^{\infty} E_r(z) dx dy, \quad (4.2)$$

allowing us to conveniently write an overlap integral between two different, normalised modes as

$$\langle E_{r1}|E_{r2}\rangle = \iint_{-\infty}^{\infty} E_{r1}^* E_{r2} dx dy. \quad (4.3)$$

Next, with the aid of figure 4.4, we define the modes involved in this analysis:

- $|\psi^\pm\rangle$ is the fibre(-guided) mode,
- $|\phi^\pm\rangle$ is the cavity mode,
- $|\psi_r\rangle = \hat{R}|\psi^+\rangle$ is the corresponding back-reflected mode,

where $+(-)$ denotes right(left) propagation and \hat{R} is a unitary transformation corresponding to the reflection. Figure 4.4 illustrates the modes, and the plane over which we take the overlap integral: the fibre/mirror interface. Now, we raise the first salient point: the single-mode fibre restricts $|\psi_r\rangle$ and the left-leaking cavity light to the same mode. Therefore, in contrast

⁷ We stress that, in this subsection only, the Dirac notation does not denote a quantum state. We use Dirac notation for these classical overlap integrals so that our analysis reads similarly to the literature from which it is derived [129, 130]. Note also in said literature and here that, unlike true Dirac notation, a bra or ket does not necessarily identify a vector quantity.

to standard free-space cavities, these two fields can fully interfere. This manifests as an over- or under-coupling of the cavity, parametrised by β : a geometry dependence which we will soon define. For example, if the wavefronts of the two fields have a different curvature (fibre field collimated, cavity-leaking field diverging), then they may interfere.

Armed with this notation, we can express all four possible overlap integrals simply as

$$\begin{aligned}
 \langle \phi^+ | \psi^+ \rangle &= \alpha, \\
 \langle \phi^- | \psi^- \rangle &= \alpha^*, \\
 \langle \phi^- | \psi_r \rangle &= \langle \phi^+ | \hat{R}^\dagger \hat{R} | \psi^+ \rangle = \langle \phi^+ | \psi^+ \rangle = \alpha, \\
 \langle \psi^- | \psi_r \rangle &= \beta,
 \end{aligned} \tag{4.4}$$

condensing them to the complex overlap amplitudes α and β . Thus, α is the mode-matched component and β is the interfering component in the fibre. In this section, we will focus only on α , but we note that the effect of β is to make the cavity resonance lineshape asymmetric: the larger β is, the greater the asymmetry. We will use this asymmetry to experimentally measure a part of the fibre geometry in section 4.4.2.

At the fibre-mirror interface, following on from the resonator mode definitions in section 2.1.2, the fibre's mode field in 1D for either left or right propagation is⁸

$$\psi_f^\pm(\rho) = \left(\frac{2}{\pi w_f^2} \right)^{\frac{1}{4}} e^{-\left(\frac{\rho}{w_f}\right)^2}, \tag{4.5}$$

⁸ Note that in [34, 129, 130], they use a different notation: they denote this scalar quantity as $|\psi_f^\pm(\rho)\rangle$ rather than $\psi_f^\pm(\rho)$. They distinguish that scalar from the vector quantity $|\psi_f^\pm\rangle$ only by including the function argument.

where w_f is the fibre mode field radius.⁹ Similarly, the cavity mode resembles equation 2.14:

$$\phi^\pm(\rho) = \left(\frac{2}{\pi w_m^2} \right)^{\frac{1}{4}} e^{-\left(\frac{\rho'}{w_m}\right)^2 - i\left(\frac{\rho'}{2z} + k\rho'\theta\right)}, \quad (4.6)$$

where w_m is the cavity mode field radius at the mirror and where we account for any misalignment between the cavity and fibre by introducing the decentered radial co-ordinate ρ' and the misalignment angle θ (in the small-angle approximation),

$$\rho' = \rho - d, \quad (4.7)$$

$$\theta = \frac{z}{\rho}, \quad (4.8)$$

where d is the decentration between the deepest point of fibre-tip mirror and the fibre core, as seen in figure 4.4. The overlap integral is a 2D integral in the xy -plane; the axial symmetry means we can express it as a product of two 1D integrals and simply replace ρ with x and y ,

$$\begin{aligned} \alpha = \langle \psi^\pm | \phi^\pm \rangle &= \iint_{-\infty}^{\infty} \psi^\pm(x, y)^* \phi^\pm(x, y) dx dy, \\ &= \int_{-\infty}^{\infty} \psi^\pm(x)^* \phi^\pm(x) dx \int_{-\infty}^{\infty} \psi^\pm(y)^* \phi^\pm(y) dy, \end{aligned} \quad (4.9)$$

where both superscripts must have the same sign (i. e. modes propagate in the same direction).

⁹ The field amplitude is normalised such that $\langle \psi_f^\pm | \psi_f^\pm \rangle = \int_{-\infty}^{\infty} (\psi_f^\pm(\rho))^* \psi_f^\pm(\rho) d\rho = 1$.

Solving the first 1D integral gives the analytical solution¹⁰

$$\int_{-\infty}^{\infty} \psi^{\pm}(x)^* \phi^{\pm}(x) dx = 2 \left(w_f w \left(\frac{ik}{\zeta} + 2 \left(\frac{1}{w_f^2} + \frac{1}{w^2} \right) \right) \right)^{-1/2} \times \exp \left(\frac{-k^2 \zeta w_f^2 w^2 \theta^2 - d^2 (4\zeta + i2kw^2) + i4k\zeta w^2 \theta d}{4\zeta (w_f^2 + w^2) + i2kw_f^2 w^2} \right). \quad (4.10)$$

Repeating this calculation for the y direction results in the 2D overlap integral, α . Then, the mode-matching efficiency is

$$\epsilon = |\alpha|^2. \quad (4.11)$$

We note that for the fibre-tip mirrors, we must add a geometry-dependent phase factor to α that depends on how much of the mirror the light must pass through before it reaches the fibre mode. Doing this, we can take into account the mirror curvature: for example, a more strongly curved mirror will have a larger phase factor at the shallower parts. Details of the phase factor are available in [129]; all we need to know is that, because we took the absolute value squared, this phase factor only affects the lineshape asymmetry and not the mode-matching. If instead we were calculating $\langle \psi^- | \psi_r \rangle = \beta$, we use a different phase factor to go from $|\phi\rangle$ to $|\psi_r\rangle$, since the light goes into the cladding.

Now, we want to compare our fibre cavity with non-fibre cavities, so that we can determine the effect of the SM fibres' mode filtering on ϵ . To prevent us from overestimating ϵ for non-fibre cavities, we must therefore remove the effect of the fibre cladding. So, we now perform an alternative calculation of ϵ for cladding-less mirrors. Removing the fibre cladding is akin to replacing the in-coupling fibre-tip with a glass substrate of the same shape but a diameter $2w_f$.¹¹ We can use the following overlap integral-derived expression for the

¹⁰ Equation 4.10 is calculated using a computer algebra system.

¹¹ An FPC made of two large glass substrates would, if the incoupling light is properly aligned, have 100 % mode-matching. This is why we consider our 'cladding-less' model as between an in-coupling glass substrate

mode-matching efficiency from this glass substrate to the cavity mode (again in the small angle approximation) [130],

$$\epsilon_{\text{glass}} = \epsilon_0 e^{-2(d/d_{\text{eff}})^2} e^{-2(\theta/\theta_{\text{eff}})^2}, \quad (4.12)$$

where ϵ_0 is the mode-matching for a perfectly centred and aligned cavity [118],

$$\epsilon_0 = \frac{4}{\left(\frac{w_f^2}{w_m^2} + \frac{w_m^2}{w_f^2}\right)^2 + \left(\frac{kw_f w_m}{2\zeta}\right)^2}, \quad (4.13)$$

and d_{eff} and θ_{eff} are the transverse misalignment and tilt angle respectively for which ϵ falls by $1/e$ [130],

$$\theta_{\text{eff}}^2 = \frac{2\lambda^2}{\pi^2 \epsilon_0 (w_m + w_f)^2}, \quad (4.14)$$

$$d_{\text{eff}}^2 = \frac{2}{\epsilon_0 \left(\frac{1}{w_m^2} + \frac{1}{w_f^2}\right)}. \quad (4.15)$$

Therefore, in figure 4.5 we now input our fibre-tip parameters from table 4.1 into both the model including the cladding and the simplified model not including the cladding,¹² thus allowing us to directly see the effect of the SM fibre mode-filtering on the mode-matching.

Examining figure 4.5, we see that the fibre cavity is more sensitive to misalignment, particularly angular misalignment, compared to non-fibre cavities (which do not have a core/cladding). This shows how the mode-filtering by the SM fibres make coupling to the fibre mode more critical: the fibre rejects any mode that is not properly aligned to it, including the well-aligned fraction of the cavity mode that interfered with the mode internally reflected

and an out-coupling fibre-tip: so that we can compare the misalignment but eliminate the mode-filtering effect from our calculations.

¹² Recall that to calculate w_m , we use equations 2.10 to 2.13.

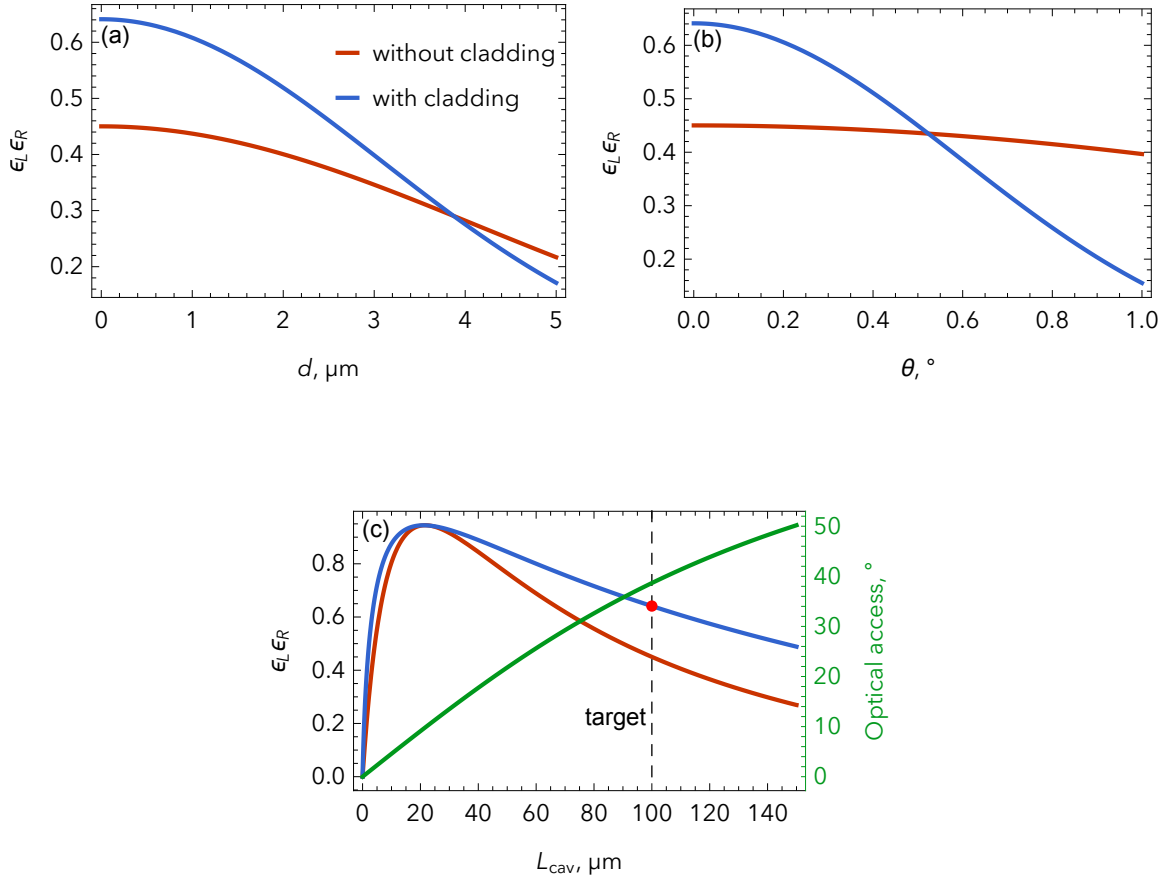


Figure 4.5: Mode-matching transmission efficiency through the cavity versus **(a)** decentration/transverse misalignment, **(b)** angular misalignment, and **(c)** cavity length and optical access. For **(a)** and **(b)**, we perform the calculation using the target length in **(c)**, $L_{\text{cav}} = 100 \mu\text{m}$. ϵ_L and ϵ_R denote the mode-matching through the left and right mirrors respectively. The two lines in each graph allow us to compare the mode-matching if we include the fraction coupled into the cladding or not. In **(c)**, we also plot the optical access, to guide us in choosing a suitable target L_{cav} , and the red dot denotes the mode-matching for our target: 0.64. For these plots, we use the parameters from table 4.1 (with $\zeta_1 = \zeta_2 = 220 \mu\text{m}$).

from the fibre. As shown in figures 4.5a and 4.5b, a transverse misalignment of $2\ \mu\text{m}$ or an angular misalignment of just 0.4° reduces the overall mode-matching by 20 %. We will use this 20 % misalignment tolerance as a benchmark for our alignment performance; although the choice itself is arbitrary, 20 % is a reasonable value to aim for given comparable loss in other cavity-based interfacing or networking schemes [21, 23].

Still to do is to pick a cavity length. Our main priority is optical access for strongly focused optical tweezers. Yet, we must strike a balance between longer cavities for tweezers access and shorter cavities for a better ϵ and g . Figure 4.5c highlights this interplay, and helps us to determine the optimal L_{cav} . Clipping loss, where w becomes larger than D_ζ and so that extraneous fraction is lost, also places an upper bound on L_{cav} ; we will discuss this in section 4.3. Therefore, to pick the cavity length we can condense this information into a single principle: make the cavity as long as possible before clipping loss occurs, and then slightly shorter than that to maintain cavity stability. Our resulting choice of cavity length is $L_{\text{cav}} = 100\ \mu\text{m}$, with a corresponding mode-matching $\epsilon_L \epsilon_R = 0.64$.

Summarising this section, we now know to expect a mode-matching $\epsilon_L \epsilon_R = 0.64$ for a well-aligned cavity with the typical values of parameters defined in tables 4.1, and also our 20 % alignment tolerance of either $2.0\ \mu\text{m}$ transverse misalignment or 0.4° angular misalignment. Moreover, by considering these mode-matching efficiencies versus optical access for tweezer light, we know to aim for $L_{\text{cav}} = 100\ \mu\text{m}$. Thus, we have a rough specification of the cavity.

However, there are some deviations from an ideal mirror that we must be aware of. In the next section, we describe how we assemble our **FFPC** to this specification.




Figure 4.6: Illustrations of various fibre-tips after ablation. In the ablation process, the final shape of the fibre-tip can deviate from the ideal shape, for reasons discussed in the main text. A non-square cleave angle will also result in the ablation afterwards being decentered by a distance d .

4.3 How accurately can we specify a fibre-tip mirror?

Earlier in section 4.1, we had noted possible manufacturing flaws in the ablation process that fall into three categories: decentration, ellipticity, and clipping loss. Figure 4.6 illustrates what each of these are. In this section, we discuss each aspect, how the resulting deviations from the ideal mirror affect the cavity performance (e. g. mode-matching), and what we do to compensate for these aspects.

4.3.1 Decentration

We have already examined the impact of decentration in figure 4.5a, but we have not discussed what causes it. Up until now, one may have implicitly assumed that the ablation laser is centred on the fibre-tip core. However, there may be a small angle between the fibre-tip and the ablation beam. Even for an angle as small as 1° , the centre of the fibre-tip depression would be off-centre by $3\ \mu\text{m}$ [131]. If the ablation is off-centre, ϵ will be reduced.

Inevitably, there will be some unavoidable decentration, because the limited sharpness of the fibre cleaver means that it may cleave the fibre at a slight angle. Also, clamping the fibre into its mount for the ablation will shift it by a slight angle. Amongst some of the fibre cavity research, this has been a significant problem [115, 129, 131, 132]. So, we must know how

much decentration we expect in our fibre-tips.

For our fibres, the decentration d of each fibre-tip was measured at the time of their manufacture using white light interferometry (WLI) [125, 126]. All we need to know is that WLI gives us information on the curvature and centre by imaging concentric bands around the deepest point of the mirror: the number of bands corresponds to ζ , in much the same way the number of rings in a tree trunk tells us its age. However, the original WLI and image data, and corresponding decentration information, is not available. Furthermore, the fibre-tip is difficult to observe under an optical microscope due to the need to bend it beyond its specified bending radius to place it in focus. At the time, we did not have a free-space optical setup suitable for looking at the fibre-tips in detail, so we instead used an atomic force microscope (AFM)¹³ to determine d by measuring the fibre-tip topography.

Briefly explaining the method, an AFM works by raster-scanning a small cantilever very close to the sample, and then converting the deflecting force felt by the cantilever (in this case, van der Waals) into a height at each point of the sample [133]. Figure 4.7 shows an AFM measure of a fibre-tip that was reserved for testing it. Because its scan range, at $(64 \times 64) \mu\text{m}$, is too small, we instead scanned from the edge of the fibre-tip inwards. Then, debris at the edge of this test fibre-tip prevented us from being able to scan across the topography of the ablation, since the cantilever has difficulty tracking the smooth, shallow ablation amongst the steep peaks of the debris. Cleaning this fibre-tip did not remove this large debris: perhaps they are chips in the fibre glass itself, as indicated by figure 4.7b. As noted in [118], these AFM measures on fibre-tips are time-consuming, and especially run the risk of breaking them while mounting them in the AFM. From our experience, we decide that AFM measurement of the topography is not the optimal way to determine d for these fibre-tips.

Having explored the AFM method, we therefore instead proceed by measuring the

¹³ Nanosurf NaniteAFM.

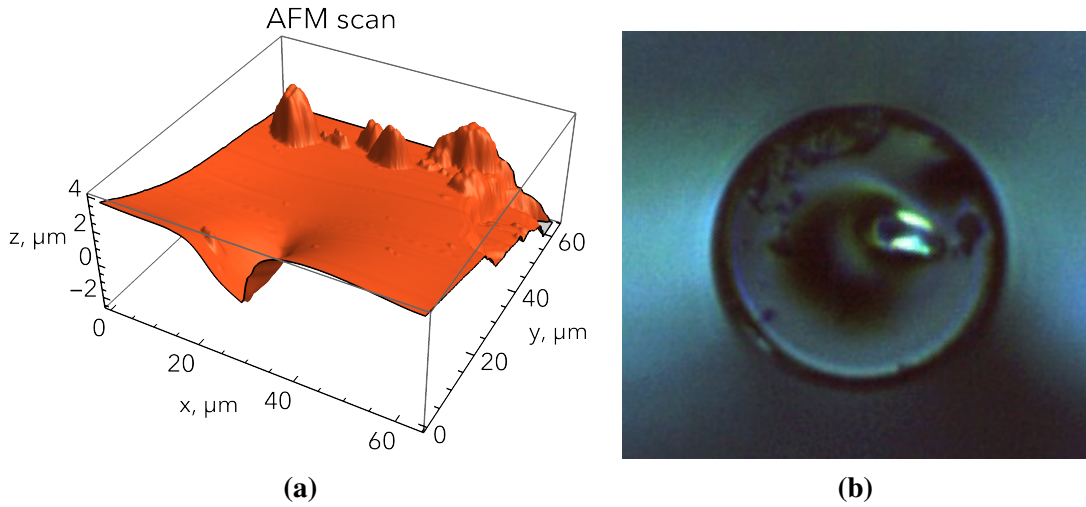


Figure 4.7: Measuring fibre-tip topography using an AFM. **(a)** AFM measurement of a fibre-tip reserved for testing purposes. A limited scan range prevents us from scanning across the full diameter of the fibre-tip ablation. There is a slight tilt in the topography, which is caused by a unavoidable tilt of the fibre-tip in the AFM mount. **(b)** Optical microscope image while under the AFM, indicating possible chipped edges of the glass.

asymmetry of the cavity resonance, and extracting d from that. For that, we must derive a link between d and the β parameter of the mode-matching analysis from section 4.2, so we defer discussion to section 4.5.

4.3.2 Ellipticity/Birefringence

Many lasers, including the CO₂ ablation laser, exhibit some ellipticity. If the ablation beam (pulse) is elliptical, then the mirror shape is elliptical too. An FPC with at least one elliptical mirror will be birefringent: that is, different polarisation eigenmodes of the cavity will have a different frequency [134]. This is because the field must vanish at the mirror surface, as mentioned earlier in section 2.1.2, yet the different polarisation modes have a slightly different wavefront curvature at the mirrors because of the ellipticity.¹⁴ Therefore, ellipticity-induced

¹⁴ This minuscule difference is manifest only for the nonparaxial case, where one solves the Helmholtz equation without the paraxial approximation. As stated in [134], birefringence due to mechanical stress of the mirror coating is often a negligible effect in comparison.

birefringence splits the cavity resonance peak accordingly.

The fundamental transverse mode, E_{00} , has two (orthogonal) polarisation eigenmodes, and so is split into two peaks separated in frequency by Δ_B . The birefringence magnitude relative to the linewidth, Δ_B/κ , can be especially large for microscale FPCs because of their ellipticity, narrow linewidths and small ζ . As a result, we must be careful: this splitting can reduce the effective g for polarised photons, because the cavity may couple to undesired atomic transitions or filter the possible polarisations that our atom-cavity system could support/output [32]. For example, a birefringent cavity might have one polarisation eigenmode couple resonantly (strongly) to the desired atomic transition, whereas the other eigenmode couples off-resonantly (weakly). Ideally, birefringence should be minimised to ensure full control of the polarisation degree of freedom for an efficient and faithful quantum interface.¹⁵

Each of our fibre-tip mirrors has a different ellipticity, which we denote using different ζ along the major and minor axes. The ellipticity is the ratio of these two values. Unlike for decentration, ellipticity data is available to us; we do not need to perform another set of WLI measurements.¹⁶ From table 4.1, the ellipticity has a standard deviation of 0.1. Thus, we pick carefully which fibre-tips to combine. Even with this large spread, if we pick two fibre-tips that have similar ellipticity, we can rotate one fibre-tip with respect to the other such that their minor and major axes coincide: this will minimise Δ_B , as we shall now see.

Figure 4.8 demonstrates our control of the FFPC birefringence when one fibre-tip is rotated and the other kept fixed.¹⁷ According to [134], the corresponding phase shift from one

¹⁵ An alternate strategy is to deliberately engineer a large birefringence into the cavity so that the two peaks are well separated, which can also increase the interfacing efficiency in special cases: for example, if the same atom is repeatedly prepared in alternating orthogonal polarisation states [102].

¹⁶ We standardise the ellipticity data collected in [125], where after doing so, the value of the ellipticity is either $\zeta_{\text{major}}/\zeta_{\text{minor}}$ or $\zeta_{\text{minor}}/\zeta_{\text{major}}$: whichever is smaller. By doing this, we place the ellipticity on the range [0, 1].

¹⁷ Using a Thorlabs HFR1. We take these measurements on our cavity characterisation setup, which we describe in detail soon in section 4.4.1.

fibre-tip is

$$\Delta\phi = \frac{\lambda}{2\pi} \frac{\zeta_{\text{major}} - \zeta_{\text{minor}}}{\zeta_{\text{major}}\zeta_{\text{minor}}}, \quad (4.16)$$

and for two fibre-tips, the round-trip phase shift is

$$\Delta\phi_{\text{rt}} = \frac{\mathcal{F}}{2\pi} \sqrt{\Delta\phi_{\text{L}}^2 + \Delta\phi_{\text{R}}^2 + 2\Delta\phi_{\text{L}}\Delta\phi_{\text{R}} \cos(2\theta_{\cup} + \theta_{\text{offset}})}, \quad (4.17)$$

where L(R) denotes the left(right) fibre-tip, θ_{\cup} is the rotation angle of one fibre-tip with respect to the other, and θ_{offset} is a fitting parameter accounting for the unknown initial orientation of the fibre-tip. From figure 4.8, we can find the value of θ_{\cup} that gives the minimal splitting, and glue the cavity in that orientation. The fit follows equation 4.17. Since the amount of splitting is mostly unique to different fibre-tip pairs (to within the error of each birefringence measurement), we can use this to identify different fibre-tips and their ellipticities. In figure 4.8, we used fibre-tips with $\zeta_{\text{major}} = 270 \mu\text{m}$ and an ellipticity of 0.87 on the left, and $\zeta_{\text{major}} = 345 \mu\text{m}$ and an ellipticity of 0.88 on the right.¹⁸ Because rotating a fibre-tip and realigning the FFPC many times is time-consuming, then for the FFPC that we will use in chapter 5, we perform an abridged version of this experiment, where we instead take only enough data points to ensure that our minimum Δ_{B} indeed lies in a valley/local minimum: 6 rather than 12 points. Therefore, we glue our assembled FFPC in the position of minimum Δ_{B} .

4.3.3 Clipping loss

We know that the cavity mode can only be supported by a spherical fibre-tip curvature (so that the field vanishes on it): thus, we need a constant ζ across the fibre-tip. However, the ablated feature has a Gaussian shape corresponding to the Gaussian CO₂ ablation laser. Are

¹⁸ The latter fibre-tip is the same as for the assembled FFPC; the former fibre-tip broke during a cavity assembly attempt later on.




Figure 4.8: By rotating one fibre-tip with respect to another fibre-tip and measuring the frequency splitting of the fundamental transverse mode Δ_B , we can find the point of minimum birefringence when aligning the cavity. Errors are the standard deviation of 8 linewidth measurements per data point. The fit follows equation 4.17. Insets show the cavity transmission intensity as we scan the cavity across resonance, where Δf is the frequency detuning from resonance; the small peaks are sidebands which act as frequency markers, and will be described soon in figure 4.11 and section 4.4.1.

the two statements compatible? Luckily, the Gaussian shape is approximately spherical in the central region of the fibre-tip. The size of this central region is what defines the effective mirror diameter D_ζ . Different conventions for the effective size exist; we use the full width at $1/e$ of the Gaussian fit in the WLI measurement to define D_ζ [125, 126].¹⁹ Since we usually have $2w_m < D_\zeta$, where $2w_m$ is typically of order $10\ \mu\text{m}$ for our experiments, then usually we accommodate all of the cavity mode at the mirror(s). However, if $2w_m \gtrsim D_\zeta$, then the cavity light falls outside the mirror region and is lost. This is clipping loss. Naturally, the longer L_{cav} is, the greater the clipping loss, because w_m grows larger with longer L_{cav} .

Using a simple model [118, 132], we can calculate what upper bound clipping loss puts on

¹⁹ We average over both the fits along the major and minor axes.

L_{cav} . Consider a simple overlap of the cavity mode at a fibre-tip,

$$\mathcal{L}_c = e^{-\frac{1}{2}\left(\frac{D_\zeta}{w_m}\right)^2}, \quad (4.18)$$

where \mathcal{L}_c is the clipping loss. For two fibre-tips, the total clipping loss would be

$$\mathcal{L}_{c,\text{tot}} = \mathcal{L}_{c,L} + \mathcal{L}_{c,R},$$

where the subscripts denote left and right fibre-tips. Now, since the finesse is, including losses [43],

$$\mathcal{F} = \frac{2\pi}{2(T + \mathcal{L}_S + \mathcal{L}_A) + \mathcal{L}_{c,\text{tot}}}, \quad (4.19)$$

with T , \mathcal{L}_S and \mathcal{L}_A defined in table 4.1, then we expect \mathcal{F} to stay virtually constant at shorter lengths, then sharply drop past the length where clipping loss starts to dominate. If we input the parameters from table 4.1, then we have a 10% drop in \mathcal{F} at $L_{\text{cav}} = 250 \mu\text{m}$. Of course, because we want better optical access and mode-matching, we would stay well below that length anyway. However, suppose one fibre-tip has $D_\zeta = 30 \mu\text{m}$: then, we instead have a 10% drop in \mathcal{F} at $L_{\text{cav}} = 140 \mu\text{m}$. For our assembled FFPC, we do find a similar drop in \mathcal{F} at $140 \mu\text{m}$, so we glue the cavity at a slightly shorter length than this to avoid clipping loss.

As an aside, we note that none of the manufacturing flaws discussed in this section are fundamental limitations of fibre cavities [132, 135]. We will discuss these fundamentals further in chapter 6.

Summarising this section, we now make sure that when we assemble our FFPC, we will:

1. Measure the asymmetry of the reflected cavity resonance lineshape in order to estimate the decentration d ;
2. Rotate one fibre-tip with respect to the other to minimise the birefringent splitting due

to the elliptical mirror shape(s);

3. Keep L_{cav} below the point where clipping loss starts to occur.

From this, we now move forwards onto the [FFPC](#) assembly.

4.4 Assembling a fibre Fabry-Pérot cavity

Now that we have set our goal to assemble a [FFPC](#) with sub-micron precision tolerance, here we demonstrate our method. Our specially developed process mitigates misalignment as far as possible throughout the assembly process: our key improvement is to not only identify, but mitigate epoxy (glue) deformation, which is the main source of misalignment in precision microcavity assemblies such as demonstrated here. The reader looking for a quick summary of our assembly method may wish to skip forward to [figure 4.15](#); how we get to and develop that method is shown before that, and how we improve on other methods is shown after.

First, we describe how to align and characterise the fibre-tips in [section 4.4.1](#). Next, we show our characterisation results prior to gluing the cavity in [section 4.4.2](#). Then—the pivotal part of this section—we describe and discuss our assembly method in [section 4.4.3](#), and follow with our mitigation and discussion of the main source of misalignment, epoxy deformation, in [section 4.4.4](#).

4.4.1 Cavity alignment/characterisation apparatus

To align and characterise the cavity in the laboratory, our setup must be precise and sensitive enough to observe transmitted and reflected light with ease while we carefully position the fibres. [Figure 4.9](#) shows our alignment setup, featuring three lasers of different wavelengths for sensitive alignment and a stable, precise translation stage for fibre positioning. The fibres are held in V-grooves, so named because a valley in a V-shape is cut into them, in which

the round fibre sits stably. Figure 4.10 shows the mirror transmission T versus wavelength, explaining why we use three different-coloured lasers: effectively, we align in three stages of ‘easy, medium, hard’, since the cavity is successively more reflective and harder to couple into at ‘633 nm, 840 nm, 780 nm’.

Initially, we position the two fibre-tips close to each other using a microscope. Then, we observe the cavity transmission when injecting light from a 633 nm fibre alignment pen²⁰. At this wavelength, R is weak enough that we can more easily observe the first transmission signal, and subsequently coarsely align the transverse positions of the fibre-tips. Once we maximise the transmission, we replace the alignment pen with a bare 840 nm laser diode and activate the cavity length scanning by driving piezos that move back and forth along the cavity axis, either on the translation stage or on the cavity mount itself. So, now we can observe cavity resonance,²¹ and align the angular position of the fibre on the stage until transmission is maximised again while minimising coupling into higher-order transverse modes. Finally, we replace the 840 nm laser diode with a more stable external cavity diode laser (ECDL) tuned to about 780.24 nm, corresponding to the ^{87}Rb D_2 line, and maximise the transmission once more. Note that although we monitor both transmission and reflection, we align on transmission rather than reflection at all stages. Maximal cavity transmission, for our symmetric cavities, will always give us maximal ϵ since we would not observe the non-mode-matched component β on transmission, as we had discussed in the previous section 4.2.

To characterise our cavity, we apply frequency modulation (FM) to the laser current, producing laser sidebands that serve as frequency markers for cavity linewidth measurements, as shown in figure 4.11. With this frequency scale, plus a measurement of L_{cav} with the microscope, we can measure \mathcal{F} using equation 2.7. At this point, we have fairly well aligned

²⁰ Laser Components FIBERCHECK

²¹ The 840 nm laser exhibits highly multi-mode oscillation, so we align on the dominant mode of the laser.


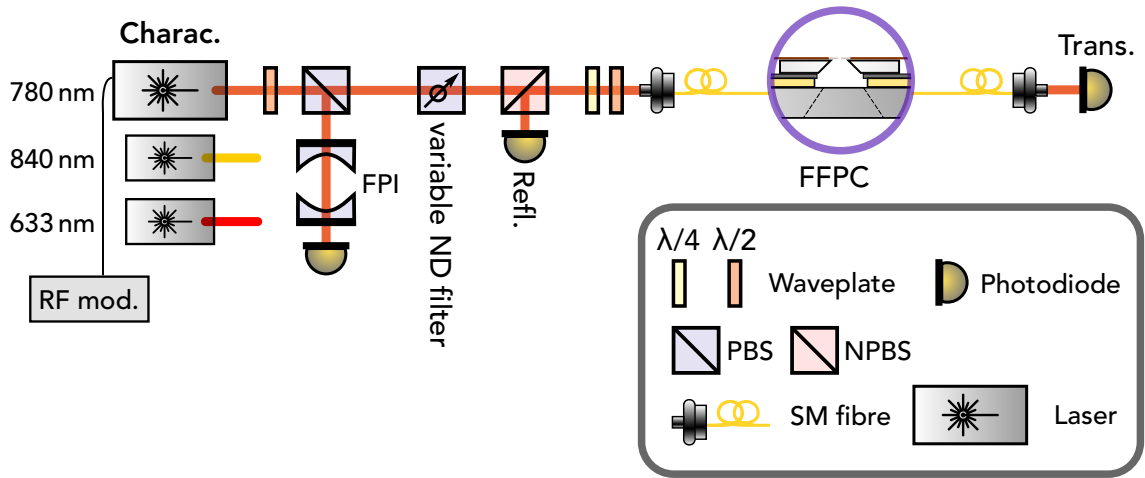


Figure 4.9: **FFPC** characterisation apparatus. **(a)** **FFPC** alignment laser setup. We send light from one of three lasers into the left fibre: the other end of the fibre is the mirror-coated one. Note that when used, the 840 nm and 633 nm lasers are connected directly to the fibre. When using the 780 nm laser, a Fabry-Pérot interferometer (FPI) allows us to check that the laser is operating on a single mode only. Should we need an additional wavelength reference, we use a fibre-connected wavemeter (not shown). We can also control the amount of input light with the variable neutral-density (ND) filter. **(b)** Photo of one permutation of the alignment setup as used to measure the birefringence of individual fibre-tips, with one fibre-tip in the fibre rotator (left) and the other in a large V-groove attached to a six-axis translation stage (right). We observe the fibre-tips from above with a microscope (inset: microscope image of aligned **FFPC**.)




Figure 4.10: How the fibre-tip mirror coating’s transmission varies with the wavelength of input light (data from the coating company ATFilms and plot adapted from [125]). The coating is optimised for 780 nm, which is where the ^{87}Rb D_2 transition lies that we will use in our experiments. We align the cavity using lasers at each successive wavelength of ‘easy, medium, hard’.

the cavity, so ϵ does not increase much with further adjustment. However, \mathcal{F} is more sensitive to the alignment, especially as L_{cav} becomes longer, because of clipping loss. Therefore, we now switch to aligning on \mathcal{F} : first, we bring the fibres close and maximise \mathcal{F} , and then we gradually bring the fibres further apart whilst making small alignment adjustments until \mathcal{F} starts to drop (i. e. we encounter clipping loss). This is how we can make long, well-aligned cavities with high \mathcal{F} and ϵ . Furthermore, this method allows us to maximise \mathcal{F} and therefore the atom-cavity coupling even if ϵ is reduced due to decentred fibre(s).

Because of the cavity’s birefringence, we must set the polarisation of the input light using quarter- and half-wave plates. Since we use **SM** and not polarisation-maintaining (PM) fibre,²² the polarisation will naturally drift over time with e.g. thermal fluctuations.

²² No **FFPCs** built from **PM** fibres appear in the literature; this may be because the internal structure of **PM** fibres, with their internal stress-inducing rods, mean they cannot be ablated in the same way as **SM** fibres.





Figure 4.11: Resonance of an FFPC, with sidebands acting as frequency markers. We use a three-peak Lorentzian fit (or six-peak, if calculating birefringence) (solid blue) to calculate $\Delta\omega_{\text{FWHM}}$. Here, $\Delta f_{\text{FWHM}} = \Delta\omega_{\text{FWHM}}/2\pi = (14.71 \pm 0.03)$ MHz.

These drifts have a characteristic timescale of a few minutes (before a waveplate needs to be adjusted), so while there are no issues with using non-PM fibres in the characterisation, we connect PM fibres to the cavity’s SM fibres for our atom-cavity experiments, where we must account for the polarisation drifts affecting the cavity resonance over a longer timescale.

We characterise and assemble FFPCs inside a flowbox, which is an optical table sealed off from the ambient environment that has a small laminar flow of clean air keeping dust down on the table floor. Should fibre-tips be contaminated, we can still recover them by cleaning them in an ultrasonic bath. One fibre-tip at a time is bathed in a beaker containing a fresh (i. e. unused) solvent. After a few minutes, we replace that beaker with another beaker containing a different solvent. We find the following solvents and bathing times to be effective:



Analytical grade reagents are used for the isopropanol and acetone, while the MilliQ water is the purest solvent available (about ten times purer than analytical reagents). With this process, we can remove organic contaminants (e. g. almost all dust from a non-synthetic

source). There are methods to clean fibre-tips with inorganic contaminants as well [136]. However, the substances used would react with the copper shroud of the fibres.²³

Our cleaning process can recover \mathcal{F} well for most mirrors: we improved \mathcal{F} for one test cavity from 3.5×10^4 to 14×10^4 after cleaning one of the fibres. Among 10–20 ultrasonic-cleaned fibre-tips, we observed no instances of reduced mirror quality.²⁴ When instead of ultrapure water we used acetone, which is more volatile, as the final solvent, and rapidly dried the substrate with clean air,²⁵ we also observed no instances of reduced mirror quality. As an alternative, we also try First Contact Polymer,²⁶ as used on the cavity mirrors in the LIGO observatory [137]. Effectively, a liquid First Contact dries on contact with the surface, and the user peels it off along with the contaminants. However, we found this to be less suitable for fibre-tips due to their small size and the strong force needed to peel the First Contact (although may work well with adjustments to the dispensation of it).

With the alignment setup completed, next comes the characterisation of the cavity that we will thereafter assemble, glue, and put into vacuum.

4.4.2 Cavity testing before vacuum

In this section, we firstly pick the best fibre-tips to make a cavity with, then secondly we test that cavity.

Choosing the best mirrors

We know that the fibre-tips exhibit a large variation in their properties. So, to make the best cavity, we need to pick the best ones. We have seen in section 4.3 how they could have

²³ See section 5.2.3 for more details about the copper shroud.

²⁴ For macroscopic mirrors ultrasonic-cleaned by other members of the author's group, we observe some instances of reduced mirror quality: this could be due to the larger surface taking longer to dry, attracting more dust or other contaminants while the final milliQ H₂O solvent dries.

²⁵ Ambersil Air Duster 2. Any jet without solvents, such as dry nitrogen, should be suitable.

²⁶ Manufactured by Photonic Cleaning Technologies.

sub-optimal shapes, be contaminated with dust, or their mirror coating could be chipped or or improperly deposited. To this end, we tested each fibre-tip against a reference mirror, as we shall now describe.

To test a fibre-tip, we firstly place it onto the translation stage and align it against a large superpolished mirror with the same mirror coating (i.e. same R), but with a much larger $\zeta = 5$ cm and $D = 7.75$ mm. Compared to the strongly curved fibre-tips, this large mirror is effectively flat. Therefore, we use this large, flat mirror as a reference. We then place the reference in a gimbal mount and, using a telescope and beam splitter which both face the non-coated surface of the reference, we capture the cavity light on both a camera and photodiode. Then, we align the cavity, and see how long we can make this cavity before \mathcal{F} starts to drop, and what the maximum \mathcal{F} for that fibre is. In both cases, the larger, the better, since it means lower losses/stronger atom-cavity coupling.

The result of our testing for many fibre-tips is displayed in figure 4.12. Although from this we now have a better idea of which fibres work well, we do have to unpick from this data a few factors that that may obscure our decision of which fibre-tips to use in our FFPC. The main factor is the systematic error in measuring L_{cav} : we are measuring it using a microscope, but the large reference mirror blocks our view of the fibre-tip. Using $\zeta = 5$ cm and $D = 7.75$ mm, we calculate the depth of the mirror—its geometrical sagitta—as $\zeta - \sqrt{\zeta^2 - (D/2)^2} = 0.15$ mm. We align the fibre-tip at the centre of the reference; given that we lose sight of the fibre-tip 150 μm before the fibre-tip could crash into the reference and break, we only push the fibre-tip in by 120 μm past this point, giving us 30 μm of leeway. However, this means our microscope measurements of L_{cav} could also have a systematic error larger than the typical 3 μm precision we can normally achieve with the microscope.

At this point, we look for a more accurate length measurement, so we instead try to measure the frequency spacing between resonances, both for the free-spectral range (FSR)




Figure 4.12: \mathcal{F} of individual fibre-tips versus L for cavities formed of one fibre-tip and one large mirror. Each colour corresponds to a particular fibre-tip (labelled as they were from the production batch). The fits, which follow the clipping loss model (see section 4.3.3), act as a guide to the eye.

and for higher-order modes. Because the cavity resonance frequencies ω_c are so well defined, we would get a more accurate measurement. However, we cannot scan our ECDL laser’s frequency ω_L far enough to capture many modes at once without the laser changing frequency discontinuously (‘mode-hopping’). Scanning instead the length with a piezo is not ideal either, since our translation stage’s piezo response is highly nonlinear if we scan over an FSR, and otherwise using the higher-order mode spacing from equation 2.18 is difficult due to the small fraction amount of light transmitted for those modes for this particular setup. Since we have enough information to see which fibre-tips work better than others, we proceed by picking first those with higher \mathcal{F} and similar ζ , and making cavities from those. Then, if a fibre-tip does not perform as well, we switch for the next best fibre-tip, and so on.

From the previous statement and figure 4.12, we therefore firstly make a cavity from fibre-tips B4 and B9 at $L_{\text{cav}} = 100 \mu\text{m}$, which gives us the optical access that we aimed for, while stopping just before the cavity starts to succumb to clipping loss. We therefore meet our target L_{cav} that we set in section 4.2.

However, fibres B4 and B9 were not used in the final assembled FFPC, for reasons we will

discuss in section 4.5. Eventually, we assembled a cavity with two similar fibres denoted B24 and A8, which have a shallower curvature than our target. Instead of $\zeta \approx 200 \mu\text{m}$ for all the fibre-tips in figure 4.12, B24 and A8 have $\zeta \approx 300 \mu\text{m}$. For the rest of this chapter, the FFPC data we show will be only for fibre-tips B24 and A8.

Testing the cavity before gluing it

We must ascertain that our FFPC will have enough optical access for the tweezers. Therefore, we see how long we can make the cavity. Figure 4.13a shows its finesse versus length. As a marker of how well-aligned the cavity must be, we need \mathcal{L} to stay roughly constant until at least our target length of $L_{\text{cav}} = 100 \mu\text{m}$, which the data indeed shows, to within error. We can actually make the cavity longer than we need, up to $140 \mu\text{m}$ without a drop in \mathcal{F} (clipping loss). However, ϵ then becomes too low to make the cavity usable for typical experiments. So, we will generally set L_{cav} close to our target length.

For our atom-cavity experiment, we also need to bend the fibres tightly. The compactness of the vacuum cell (see section 5.2.3 for details) means that, to fit the FFPC, we must bend one fibre with a minimum bend radius of 0.5 cm. Normally, such small bend radii will cause large losses along that fibre: prior to placement in vacuum, these fibres were tested with gradually smaller bend radii to determine the amount of loss. The data sheet for these fibres²⁷ state a minimum long-term bend radius of 2.0 cm, but does not state the amount of loss: presumably, the tolerance for loss as stated by the manufacturer is fairly strict. We want to check what loss to expect, so we measure the transmission loss through a single flat, uncoated fibre-tip focused onto a photodiode as we bend it. Figure 4.13b shows only a small loss of light up to 0.5 cm from these particular copper fibres. Therefore, we determine bending loss not to be an issue requiring action in our setup. Regardless of the implementation, this is not

²⁷ Oxford Electronics SM800-125C. Note that the company is now defunct, so they are not able to clarify the data sheet.




Figure 4.13: Pre-vacuum cavity testing. **(a)** How long the FFPC that is eventually implemented in the vacuum chamber can be made before resonator losses start to increase. \mathcal{F} stays constant to within error throughout. But, after $140 \mu\text{m}$, the transmitted power/overall mode-matching is slightly too small to reliably capture the small sidebands needed to calculate this data. Errors are the standard deviation of 8 linewidth measurements per data point. The fit follows the same clipping loss model from equation 4.19. We plot this fit at lengths past where the data lies to highlight that our cavity characterisation lies well outside the clipping region. **(b)** Transmission loss through a single fibre as it is bent in a semicircle of gradually tighter radius. We normalise the data such that the highest measured transmission is unity. Error bars correspond to the base level of noise of the photodiode used.

a fundamental limit on FFPCs loss in general [138], because we can of course use a larger vacuum cell: here, we want to primarily make sure that optical access is fine.

Most importantly, we must measure the mode-matching before we glue the cavity, since we want to know how precise the alignment with our assembly method is. We do this by aligning the cavity as well as possible, measuring the transmission and reflection simultaneously, and then using the model from Hood et al. [34] to convert our measurements to the overall mode-matching value $\epsilon = \epsilon_L \epsilon_R$.

From equation 4.19, if we align the cavity at our target length, then we are outside the clipping loss region and equation 4.19 reduces to

$$\mathcal{F} = \frac{\pi}{T + \mathcal{L}}, \quad (4.20)$$

where $\mathcal{L} = \mathcal{L}_S + \mathcal{L}_A$. We have assumed that both fibre-tips have similar properties since they came from the same coating run: that is, $T = T_L = T_R$ and $\mathcal{L} = \mathcal{L}_L = \mathcal{L}_R$. Next, we set equation 2.4 on resonance to get

$$\frac{P_t}{\epsilon P_{\text{in}}} = \frac{T}{T + \mathcal{L}}, \quad (4.21)$$

where P_t and P_{in} are the transmitted and input powers. Note that only the mode-matched fraction of the input ϵP_{in} will couple into the cavity. Combining equations 4.20 and 4.21 then gives

$$\frac{P_t}{\epsilon P_{\text{in}}} = \left(\frac{T\mathcal{F}}{\pi} \right)^2. \quad (4.22)$$

Likewise, consider also the reflected power P_r . As we discussed in section 2.1.1, on resonance, P_t increases and P_r decreases by \mathcal{F}/π compared to their off-resonant values. Including losses and only considering the mode-matched component that contributes to the P_r that reaches our reflection photodiode, we derive a similar expression for our reflection measurement,²⁸

$$\frac{P_r - (1 - \epsilon_L)P_{\text{in}}}{\epsilon P_{\text{in}}} = \left(\frac{\mathcal{L}\mathcal{F}}{\pi} \right)^2. \quad (4.23)$$

Since that the two fibre-tips have similar geometries, for this particular calculation we make the approximation $\epsilon_L \approx \epsilon_R \approx \sqrt{\epsilon}$. Finally, we combine equations 4.22 and 4.23 to get the following two expressions that fully specify T , \mathcal{L} and ϵ :

$$\frac{P_t}{P_r - P_{\text{in}}} = \frac{\left(\frac{T\mathcal{F}}{\pi} \right)^2}{\left(\frac{\mathcal{L}\mathcal{F}}{\pi} \right)^2 - \frac{1}{\sqrt{\epsilon}}}, \quad (4.24)$$

$$\epsilon = \frac{P_t}{P_{\text{in}}} \left(\frac{\pi}{T\mathcal{F}} \right)^2. \quad (4.25)$$

In figure 4.14a is the data we use to calculate these quantities. For that data, which

²⁸ Where the incoupling fibre is the left fibre, hence we have the term with ϵ_L .

was taken immediately before the cavity was glued, we have $L_{\text{cav}} = (90 \pm 2) \mu\text{m}$, and $\mathcal{F} = (88 \pm 3) \times 10^3$.²⁹ We can see that our FFPC is well aligned here, since we do not observe any asymmetry in the reflection trace: recall from section 4.2 that the more misaligned the FFPC, the more asymmetric the reflection trace.³⁰ Inputting \mathcal{F} into equation 4.20, we measure $T + \mathcal{L} = 36 \pm 1$ ppm. Simultaneously, we measure $(P_t, P_r, P_{\text{in}}) = (1.80, 40.3, 44.0) \mu\text{W}$. Therefore, we get $T = 23 \pm 1$ ppm, $\mathcal{L} = 12 \pm 1$ ppm, and finally $\epsilon = 0.35 \pm 0.02$.

Even though the geometry of this cavity is slightly different (shallower curvatures) compared to our calculations in figure 4.5, our value of ϵ does seem below our target slightly. For those calculations with $\zeta = 220 \mu\text{m}$, we had $\epsilon = 0.64$. For this cavity, repeating the calculation with the adjusted geometry gives $\epsilon = 0.52$. Given this, we postulate the reason for lower than expected ϵ is decentration of these fibre-tips: if we assume each fibre-tip has a not unreasonable $d = 2.5 \mu\text{m}$, then we get $\epsilon = 0.34$, which agrees with our measured ϵ . As we shall now see, decentration may indeed be the cause.

Expanding on our comment that we do not observe any asymmetry in figure 4.14a, we had stated in section 4.3.1 that we will draw a link between the asymmetry and the decentration d of a fibre-tip. Now, we do that. The magnitude of the asymmetry depends on the geometry and alignment of the fibre cavity; thus, it depends on the overlap integrals α and β from section 4.2. By considering the fraction coupled into the fibre mode and into the cladding, Gallego et al. derive the following relation for the asymmetry factor [129]:

$$\mathcal{A} = \frac{\text{Im} \left[\beta \cdot (\alpha^2)^* \right]}{\text{Re} \left[\beta \cdot (\alpha^2)^* \right] - \frac{\epsilon_L^2 T}{2\mathcal{L}}}. \quad (4.26)$$

We have an analytical form for α from equation 4.10. We also have an analytical form for

²⁹ Averaged over 16 linewidth measurements, similarly to for other finesse data we presented.

³⁰ Minimal asymmetry ($\beta \approx 0$) is also a requirement for usage of equations 4.24 and 4.25, since otherwise the depth of the reflection dip changes because of the cavity being over/undercoupled [129].




Figure 4.14: Mode-matching and asymmetry measurements of the final cavity shortly before it is glued. **(a)** Simultaneous transmission (red) and reflection (blue) measurement of our well-aligned FFPC. The dashed yellow line labels P_{in} . Note we have already subtracted the background photodiode levels for both traces here. **(b)** Asymmetry measurement for a cavity where we slightly tilt the incoupling fibre. The blue lines show both the asymmetric fit of equation 4.27 (solid) and a typical symmetric Lorentzian fit (dashed) for comparison.

β , which is similar to α other than a different geometry-dependent phase factor since the light goes into the cladding, and a different beam waist since we consider the overlap with the non-guided (cladding) modes.³¹ Furthermore, Gallego et al. also derive the asymmetric Lorentzian expression that allows us to determine the asymmetry by measuring it directly,

$$\frac{P_t(\Delta f)}{P_{in}} = -A_{\text{dip}} \left(\frac{1}{1 + \Delta f^2} - \mathcal{A}_{\text{meas}} \frac{\Delta f}{1 + \Delta f^2} \right), \quad (4.27)$$

where Δf is a (normalised) detuning from resonance and A_{dip} is the amplitude of the reflection dip. Therefore, we can fit this to our reflection data and immediately obtain a value for the experimentally measured asymmetry, $\mathcal{A}_{\text{meas}}$. All that remains is to find the values of angular misalignment and decentration θ and d that best fit $\mathcal{A}_{\text{meas}} = \mathcal{A}$ in equation 4.26.

We perform this fit to the data in figure 4.14b. Compared to figure 4.14a for the well-aligned cavity, we tilt the incoupling fibre by about 1° on the alignment stage. This results in the

³¹ For clarity, we do not reproduce the expression here: it can be found in the supplemental material of [129].

expected asymmetric reflection lineshape. Then, we calculate $\min(\mathcal{A}(\theta, \phi, d) - \mathcal{A}_{\text{meas}})$, where \mathcal{A} is a function of θ , d and an arbitrary initial phase ϕ (just as α and β in equation 4.26 are), and $\mathcal{A}_{\text{meas}}$ is our experimentally measured value. We find the best fitting value is $d = 2.5 \mu\text{m}$. If we additionally leave θ as a free parameter, we also find $\theta = 1.0^\circ$, which agrees with how much we tilted the fibre, further showing that our cavity in figure 4.14a is well-aligned.

The above asymmetry measure gives $d = 2.5 \mu\text{m}$ for one fibre-tip. Although we were not able to perform the asymmetry measurement for the other fibre-tip since the cavity was being assembled at this time and it was already glued to the assemblage, it is reasonable to assume the decentrations are similar. With this assumption, then indeed our theory and measured values of ϵ agree to within experimental error.

Having performed the pre-vacuum testing, now we present our method of assembly.

4.4.3 Cavity mount assembly

Once the cavity is aligned, we must keep it aligned. As we discussed at the beginning of this chapter, we keep the cavity permanently aligned rather than actively align it in vacuum. To do this, we glue the fibres to a monolithic structure, permanently fixing their position. We call this structure the cavity mount. Our attention is therefore on ensuring that we incur as little misalignment as possible, given that the glue could deform and the structure could only be machined to a particular tolerance. How do we ensure the cavity is permanently well-aligned?

To address these issues, here we develop an assembly procedure for the [FFPC](#) with sub-micron alignment precision. Qualitatively, we improve on the alignment performance of similar structures both actively [139] and passively [131, 140, 141] aligned, where those works referenced all report some misalignment in vacuum and for their [SM-MM](#) cavities:




Figure 4.15: How we assemble the FFPC. We describe each step in the main text. An isometric perspective of the cavity mount is shown in figure 4.16.

we observe only a small drop of ϵ of 0.06 for our **SM-SM** cavity, as we will describe in section 4.5.2. Our **FFPC** assembly method is illustrated in figure 4.15, and is as follows:

1. Place the cavity mount (shown in figure 4.16) on a stable platform next to a translation stage.³²
2. Glue two shear piezos³³ and one quartz spacer to the mount using thermal-curing epoxy.³⁴
3. Glue one glass V-groove to the spacer using ultraviolet (UV)-curing epoxy.³⁵ As figure 4.17 shows, we position its front edge 250 μm from the centre of the coned aperture of the mount and parallel to the mount sides, for best optical access.
4. Glue one fibre-tip into the V-groove using UV epoxy. We position the fibre-tip with the translation stage such that the cavity waist will be at the centre of the cavity mount's aperture. Crucially for cavity stability, we also control the fibre-tip's overhang from the V-groove with micrometre precision. Here, we use a 200 μm overhang. We then thermally bake the assembly again to prevent unwanted shrinkage of the epoxy during vacuum bakeout later on.
5. Rotate the half-assembled cavity mount 180°, then align and rotate the second fibre-tip to form a minimally birefringent cavity.
6. Keeping its orientation, glue the second fibre-tip into a V-groove, also with a 200 μm overhang.
7. Turn the mount upside down and fix it to a separate platform. Then, align the fibres

³² Thorlabs MAX602/M

³³ Noliac CSAP02

³⁴ EPO-TEK 353ND

³⁵ EPO-TEK OG198-54



Figure 4.16: Cavity mount used to assemble the [FFPC](#), showing the conical aperture to allow optical access for the focused tweezer beam. Recessed regions guide the shear piezos into the correct position. Screw holes at the corners allow us to implant it into the atom-cavity experiment; we discuss its implementation in section [5.2.1](#).

again. Measure the gap, δ , between the second V-groove and the half-assembled mount.

We measure δ using the translation stage and a microscope with an accuracy of $10\ \mu\text{m}$.

8. Polish a second spacer to a thickness of $\Delta = \delta - 50\ \mu\text{m}$ and glue to the mount using thermal-curing epoxy. By doing this, we therefore ensure that the thickness of the final epoxy layer will be $50\ \mu\text{m}$, to a precision of $10\ \mu\text{m}$.³⁶ We use a sharp implement to keep the spacer sides parallel to the mount.
9. Glue the second V-groove to the second spacer using [UV](#)-curing epoxy. We aim the [UV](#) light at the sides of the cavity firstly (into and out of the plane as on the figure), and then remove the mount from the translation stage and platform and illuminate from the top (bottom as on the figure) to ensure that the epoxy fully cures.
10. The cavity mount assembly is complete.

With our method, we assemble our [FFPC](#) in a compact, monolithic device: $(35 \times 13 \times 5)\ \text{mm}$ in size, or about half the size of a pack of chewing gum. Emphasising its small size, figure [4.18](#) shows our assembled cavity before we put it in vacuum. As we mentioned at the beginning

³⁶ Δ can be controlled to an accuracy of $10\ \mu\text{m}$ with our polishing, therefore allowing control of the thickness of the [UV](#) epoxy layer to a similar precision. Together with the error of δ , we have the final layer precision $\sqrt{10^2 + 10^2} = 10\ \mu\text{m}$ to 1 s.f.




Figure 4.17: Microscope image of V-groove alignment on the cavity mount. We use a 3D-printed holder for the V-groove (not shown, right side of picture), plus a sharp implement, to position the V-groove.

of this chapter, a compact device affords us better optical access, and potentially a more mechanically stable cavity, since the fibres are fixed to the same mount rather than to two separate mounts which could vibrate independently. In practice, the cavity mount is a small part of a larger experimental system, and as such is still susceptible to vibrations if another part is not stable: in chapter 5, we will detail an external problem that affects the cavity stability.³⁷

Most importantly, our assembly method allows us to permanently align the **FFPC** with sub-micron precision, as we will show in section 4.5. How do we achieve this? As implied in the descriptions of the assembly steps, we have achieved this by using the minimum of glue in the assembly by precisely polishing the spacers to the desired thickness, and by doing the final gluing step upside down so that gravity does not pull down a fibre and misalign the cavity when we cure the **UV** glue. Epoxy deformation and gravity pulling down the fibres has significantly affected other permanently aligned **FFPCs** [131, 139, 140],³⁸ so our method

³⁷ With regards to this stability, we did successfully construct an identical mount made of (electrically insulating) Macor rather than (conducting) steel, although not with our final assembled **FFPC** due to time constraints. Such a mount would allow us to operate each of the two piezos independently: one for scanning the cavity, and one for stabilising the length on resonance (‘locking’, see section 5.3.2), which require large and small voltages respectively.

³⁸ The effect of gravity pulling down the fibres during the cure was observed in [139].




Figure 4.18: Assembled **FFPC**, before we put in into vacuum. The fibre-tips are glued to the glass V-grooves. Below those are quartz spacers. The beige ceramic plates are the shear piezos, with Kapton-insulated electrical wires attached above and below.

is a significant improvement.

Bearing in mind that the fibre diameter is $125\ \mu\text{m}$ and yet the spacer gap is only $(50 \pm 10)\ \mu\text{m}$, we have therefore left a gap small enough to still be able to align it even if a fibre is decentred or tilted. Since the rest of this structure is fixed, then effectively the only limit to how precisely we can permanently align these cavities is how small we can make this spacer gap. Our polishing process could probably be achieved with a precision of a few microns, and we could dispense small quantities of glue using micropipettes. So, if the fibre-tips are manufactured precisely (well-centred, cleaved perfectly square), then we could reduce the spacer gap also to a few microns. Thus, we can expect to surpass even our sub-micron precision with these improvements to our process.

Additionally, in the future we could further improve on our assembly method by using two translation stages: one for each fibre-tip. In this way, we can simplify the method by gluing everything at once while the mount is upside down, thereby reducing the number of steps. More pertinently, we would both make our method safer and also be able to automate it: one disadvantage of our current method is that it is time-consuming and, since we perform many

manual steps, prone to damaging fibre-tips by human error. The fibre-tips, which are glass and as thin as a human hair, are fragile. With two translation stages, we can automate the assembly once the fibre-tips are brought close, and therefore eliminate any errors that could occur if we assembled manually.

We have mentioned how minimising epoxy deformation is crucial to keeping the cavity well-aligned: in the next subsection, we highlight why this is the case.

4.4.4 Counteracting misalignment caused by epoxy deformation

Once assembled, we implant our [FFPC](#) into our atom-cavity experiment, in which it is subject to harsh conditions. It must be able to withstand extremely low pressures in ultrahigh vacuum (UHV) of $<10^{-7}$ Pa, and bakeout in vacuum of at least 100 °C for 2 h, and at most 180 °C for 2 weeks.³⁹ This requires that we minimise any outgassing—the evaporation of trapped gases—and that the cavity does not deform during bakeout and become misaligned. To meet these requirements, we use epoxy resin adhesive to glue the cavity assembly.⁴⁰ We did not use a fully mechanical assembly, i. e. fixing the cavity alignment using metal clamps, because it both adds bulk and is difficult to engineer to the required tolerance. However, the way epoxy works means that the harsh conditions may still cause it to deform irreversibly. How can we manage the resulting misalignment so that the cavity stays aligned in vacuum?

As a first step towards countering this misalignment, we must first understand the mechanism of deformation. Epoxy adhesives are polymers with either one or two chemical constituents: long monomers containing epoxide functional groups, and short amines, as both shown in [figure 4.19a](#). The long monomers are relatively viscous, so we need the structure to set to create the strong adhesive bond. When the monomers and amines are mixed together and provided with energy in the form of heat or [UV light](#), this triggers a chemical cross-linking

³⁹ Above 180 °C, the vacuum cell (see [chapter 5](#)) may fail.

⁴⁰ Note that we will only consider low-outgassing epoxy in the following and for our experiments.

reaction between the epoxide and amine functional groups, binding the monomers together into a hard, rigid 3D polymer with vastly increased mechanical and adhesive strength. This chemical cross-linking reaction is known as curing the epoxy.


Naturally, the higher the density of cross-links, the stronger and better-adhering the epoxy is: cross-linking pulls the structure together. As a result, when cured, the epoxy shrinks. So long as the assembly parts are held in position during the cure, this shrinkage will not misalign the cavity because the resulting force is fairly weak. However, this situation is complicated by two factors: the increased mobility of the polymer chains as they cure (with heat especially), and the small fraction of the epoxy mixture that does not initially cure.⁴¹ Therefore, the cure causes the polymer chains to unwind slightly. The higher the temperature, the greater the unwinding effect. At the glass transition temperature, τ_g , the epoxy changes from a glass-like structure to a rubber-like structure and so the polymer chains are able to unwind past their cross-linked sites, suddenly becoming more mobile. Once the epoxy cools down again below τ_g , the monomers remain partially unwound, hence the inelastic deformation [142]. We postulate that this was the reason for the bakeout-induced misalignment in the direction of gravity observed by Ballance [139].

With this in mind, we can condense our above discussion into two simple rules when we choose and use epoxy:

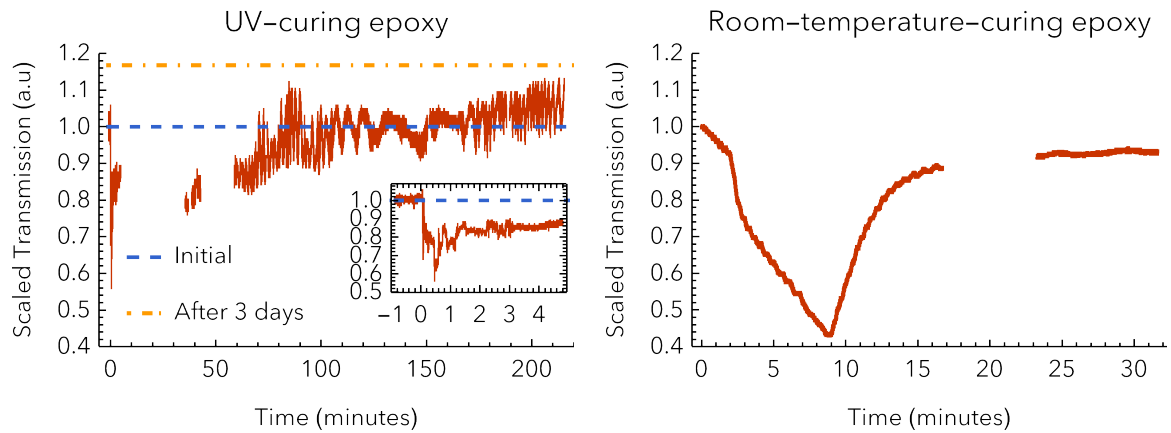
- Choose an epoxy with as high a glass transition temperature, τ_g , as possible to reduce the deformation upon bakeout.
- Cure the epoxy at room temperature, so that the **FFPC** stays aligned at its operating temperature.

For our cavities, we must know by how much the cavity is misaligned by epoxy deformation. At the same time, our fibre-tips are a finite and precious resource, so we cannot afford to

⁴¹ How small this fraction is depends on the mobility of the chains and number of available cross-linking sites.



(a)



(b)

Figure 4.19: Epoxy curing mechanism and bakeout simulation tests. **(a)** When either one or two molecule types—an epoxy, and an amine—are provided with energy either from heat or from **UV** light to start a chemical reaction, they form a bond. Our simplified sketch highlights only the functional groups, so that the resulting cross-linked structure of the cured epoxy resin is clear. In the wound case, the structure is slightly smaller, and therefore the epoxy shrinks. **(b)** Simulated bakeouts. The **UV**-curing Epo-tek OG198-54 underwent a bakeout simulation from time $t = 0$ for around 4 h. The first unbroken region (also shown in inset) is when the initial **UV** cure from the top of the ‘cavity’ took place at time $t = 0$. The second region shows a further **UV** cure from the sides. The third unbroken region is when we baked it. We observed no fibre movements on the microscope. For the room-temperature heat-curing Epo-tek 302-3M, we gradually heated the test mount to 35°C after 8 minutes, after which time we stopped the bake because the transmission dropped by too much to make this epoxy suitable for cavity assembly.

spare a cavity solely to determine the misalignment from epoxy deformation. Therefore, we quantify the misalignment by using two flat, uncoated fibre-tips made by cleaving acrylate SM fibre.⁴² We align and glue these two acrylate fibre-tips on an identical cavity mount to the one we use in the atom-cavity experiments. We forgo the shear piezos, and use two spacers of different thickness: 1.00 mm for the left fibre-tip and 1.20 mm for the right. Then, we fill the resulting 200 μm gap with epoxy: note that this gap is up to four times thicker than we use in the final FFPC. We repeat this assembly for all the different kinds of glue we want to test. Then, we simulate bakeout placing a heat lamp in close proximity to the ‘cavity’ while measuring the transmission, and also the temperature with an infrared (IR) thermometer.

Figure 4.19b shows examples of such tests. For the OG198-54 UV epoxy, we postulate the transmission increase is due to the epoxy unwinding during the bake; the oscillations may have been due to convection currents in the hot air surrounding the epoxy. Three days later, once the assembly has cooled down, we found the final transmission value had increased to 117 % of the initial value. Therefore, we register this as an 17 % loss, assuming the fibre was initially slightly misaligned and then moved to a more optimal position and that the laser power was stable. Converting this into an expected misalignment using our mode-matching theory in section 4.2, this corresponds to a transverse misalignment of 1.9 μm .⁴³ Table 4.2 shows the bakeout performance of all the epoxies we tried, where we see that Epo-tek OG198-54 performs better than the others due to its high τ_g and UV curing mechanism. The other epoxies withstand much less heat for a similar transmission loss. Therefore, we use Epo-tek OG198-54 epoxy in the final FFPC for our atom-cavity experiments.

One issue on the repeatability of these tests is how to control the amount/thickness of epoxy used. We use the sharp tip of a syringe needle without a plunger to place many beads

⁴² Oxford Electronics SM800-125P

⁴³ Note that this is below the accuracy of our microscope. Using it, we did not observe any fibre position changes, either in length or transversely.

Epoxy	Cure type	τ_g ($^{\circ}\text{C}$)	$\tau_{\text{bake,max}}$ ($^{\circ}\text{C}$)	t_{bake} (minutes)	T_{loss} (%)
Epo-tek 302-3M	Heat	55	35	8	8
Dymax OP-67-LS	UV	86	70	60	12
Epo-tek OG198-54	UV	131	85	120	17

Table 4.2: Comparison of epoxies subjected to simulated bakeouts. For the first two epoxies, we either reach a lower maximum bake temperature, $\tau_{\text{bake,max}}$, or bake at that temperature for a shorter time, t_{bake} , since by that time we will know if the epoxy deforms too much for use in vacuum. Here, T_{loss} is transmission loss.

of epoxy one at a time, since more precise microlitre syringes are expensive. In any case, since we use much less glue in the final assembled [FFPC](#) than we did in the test, and therefore we expect the final transverse misalignment to be well below $1.9\ \mu\text{m}$, then this method of dispensation is sufficient for the final assembly. In future, we can further investigate the epoxy deformation effect by using one curved fibre-tip mirror against a flat micromirror coated with small fluorescent nanoparticles that can be used as a much more accurate reference point for the position [143]. Nevertheless, our epoxy tests are time-consuming and we now have a good idea of what to expect with in our final assembly process.

To summarise this section, we have described and discussed the cavity characterisation and assembly method for our [FFPC](#). Now that we have the tools, we will use them to examine how well our final assembled cavity performs: we report the results of our assembly in the next section.

4.5 Putting the cavity into vacuum

Once we have aligned the [FFPC](#) and cleaned the cavity mount, the cavity is ready to go into [UHV](#). Recall our aim to closely watch any degradations in performance during this process, especially mode-matching. Here, we describe two attempts to put the cavity in vacuum, and how we improved on our process for our final [FFPC](#).

4.5.1 First attempt

Our first attempt to put an [FFPC](#) in vacuum was ultimately unsuccessful.⁴⁴ We observed cavity transmission when immediately testing it after closing the vacuum, but within the next few days the transmission rapidly dropped off until we observed no transmission. Therefore, we proceed by diagnosing what caused the failure, and then suggest fixes to ensure the success of the next attempt. Let us diagnose each issue in turn:

- \mathcal{F} fell from 6.9×10^4 to 3.5×10^4 . Diagnosis: either the cavity mirrors picked up contamination during implantation, or we may have observed a higher-order cavity mode rather than the fundamental one. In the latter case, the mode primarily forms on a different area on the mirror; perhaps it may already have been contaminated.
- One of the two shear piezos broke in vacuum. Diagnosis: the wires are made of Kapton, which exhibits a shape memory. They were not taped down during assembly, and so the tension built up when bending the wires through the tight vacuum chamber eventually was released and snapped the piezo.⁴⁵
- Mechanical vibrations made the cavity unstable; we were not able to attempt cavity locking⁴⁶ before the cavity failed. Diagnosis: the overhang of the fibres from the V-grooves was 3 mm rather than 250 μm for this cavity.
- We could not scan an [FSR](#). Diagnosis: with only one working piezo, although the free stroke length⁴⁷ of the piezos used is 1.5 μm and an [FSR](#) here is 0.390 μm , the epoxy

⁴⁴ We made this particular cavity with the fibre-tips labelled B2 and B16. Even earlier, we also made a cavity with fibre-tips B4 and B9, using thermal-curing epoxy only. As we found out, when the assembly cooled down after the high heat (150 °C), the cavity misaligned by $\sim 50 \mu\text{m}$. Therefore, microcavities should not be assembled using only thermal-curing epoxy.

⁴⁵ The optical fibres themselves, having a copper shroud, exhibit this memory effect more strongly and therefore are prone to break in the same fashion.

⁴⁶ See chapter 5 for details on cavity locking.

⁴⁷ Free stroke length is how far the piezo moves when driven from its maximum negative voltage to its maximum positive voltage. For our piezos, the corresponding voltage is $\pm 320 \text{ V}$; our amplifiers can output $\pm 175 \text{ V}$.

may have cured improperly, restricting the motion of the piezo.

- The UV epoxy did not fully cure. Diagnosis: our UV light source was not powerful or focused enough to cure it. We detail this issue in the main body.

Figure 4.20 shows what occurred after a few months in vacuum: over time, the cavity transmitted less and less light until resonance was no longer able to be observed. The epoxy had deformed. We used a homemade UV LED device consisting of 4 surface mounted light-emitting diodes (LEDs)⁴⁸ in a series circuit on a small breadboard. We brought them within a few cm of the V-groove to be cured. Because the LEDs exhibited a wider than expected manufacturing variation, driving them in series meant we had less control over the voltage drops across each LED. Thus, we could drive them only with enough power to cure most, but not all, of the epoxy. Since the epoxy-adhered joint was solid, we did not notice this partial cure until the cavity was in vacuum. Because of this, we used a commercial UV curing system⁴⁹ for the next FFPC, successfully.

While eventually the FFPC failed to scan at all, we still had about one month of partial operation in which we could scan a single resonance only. Because we had to input a relatively large amount of light to the cavity—10 mW CW into the fibre connector—just to be able to see the resonance, we observed thermal self-locking [125, 129, 144]. Thermal self-locking occurs when the effects of the heating of the FFPC and the scanning piezo balance each other out. When resonantly driven, the cavity mirrors absorb some fraction of circulating light, heat up, and expand due to the heat, thereby shortening the cavity. However, the scanning piezo is moving forwards for half of its scan and backwards for the other half. Thereby, the cavity lengthens during the backwards stroke. These two effects compete to keep the cavity on resonance, giving many closely spaced resonant peaks as shown in figure 4.21.

⁴⁸ Nichia NCSU276AT-0365

⁴⁹ Thorlabs CS2010




Figure 4.20: Microscope images of a fibre cavity that failed in vacuum. The top and bottom images show a side and top-down perspective respectively. The cavity is visibly misaligned by $30\ \mu\text{m}$ in the top-down perspective, and as such the cavity mode is no longer able to be driven at resonance.

Curiously, not only did we observe thermal self-locking during the piezo's backward stroke, we also observed it during the forward stroke. In that case, the thermal effect and the forward stroke both shorten the cavity, so something else must be lengthening it. From this, we can deduce that mechanical vibrations are shortening the cavity. We identify the main source of mechanical vibrations here as the long fibre overhang: the fibres extend about 3 mm from the v-grooves. At the time of this first attempt, cavity designs that were then newly published had shortened the overhang to increase stability enough for frequency-locking of L_{cav} [129, 140]: therefore, we reduced the overhang also in our second attempt.

4.5.2 Second attempt

Our second attempt to put the assembled FFPC in vacuum was mostly successful, since our rectifications after the first attempt were mostly successful. Let us examine each of our improvements in turn:

- We reduced the fibre overhang from 3.0 mm to $250\ \mu\text{m}$.




Figure 4.21: Thermal self-locking effect in our first **FFPC** in vacuum. Many resonances are closely spaced together. **Left:** **FFPC** heats in the same direction as the scan, but vibrations push it back onto resonance. **Right:** **FFPC** heats in the opposite direction to the scan. Note the difference in time scales.

- We reduced the piezo width from 7.0 to 3.5 mm. This is because having wider piezos would block the tweezer beams, since the V-grooves are now closer together.
- We tapered the V-grooves with a 45° cut, as shown in figure 4.22. Again, this is because non-tapered V-grooves would block the tweezer beams.⁵⁰
- We used a commercial **UV** curing system to ensure a successful cure.
- We changed the cavity mount dimensions slightly so that the cavity mode would sit exactly in the centre of the vacuum cell (see chapter 5 for more details).

Our most important result from our second attempt is that \mathcal{F} did not measurably drop and ϵ dropped only slightly, after both the **UV** curing and the placement in vacuum.

\mathcal{F} stayed the same at $(88 \pm 3) \times 10^3$, even while exposed to air for half a day when implanting the **FFPC** into the vacuum chamber. While this value is less than would be

⁵⁰ Due to time and resource constraints, we did not taper the left and right sides of the V-grooves close to the fibre-tip. Although we do not need it for our current experiments, doing so in the future would allow optical access for a strongly focused beam from the side—for example, a perpendicular set of tweezer beams for a 3D lattice-like trap structure.




Figure 4.22: Microscope image illustrating V-groove taper for enhanced optical access.

expected given the specification from the coating company, we can confirm that our assembly process is suitable for high finesse cavities.

ϵ dropped from 0.35 ± 0.02 to 0.29 ± 0.02 . We did not notice any change at all in alignment, cavity properties etc. after the UV top-cure, nor did we at the beginning of the UV cure from the sides. However, later on during the UV cure from the sides, the author knocked the translation stage, misaligning the cavity by enough that the transmitted power was greatly reduced. By then, although the epoxy was mostly cured from the top, a sequence of gentle adjustments allowed us to recover the transmission, after which we fully cured the epoxy from the top and sides. We did not observe any change in length from $L_{\text{cav}} = (90 \pm 2) \mu\text{m}$ before, suggesting that this was a transverse misalignment only. With this assumption, we calculate using the model in section 4.2 that the drop in ϵ corresponds to a transverse misalignment of $0.7 \mu\text{m}$. Whether the alignment was completely recovered after the translation stage was knocked is difficult to ascertain. Nevertheless, we can claim sub-micron alignment accuracy with our FFPC assembly method, which meets our alignment goal.

However, two issues with the final cavity were not totally resolved, which we now discuss.

Scanning the cavity

The first issue is our difficulty in scanning the cavity. We can just about scan an FSR when driving both piezos: with one piezo only, we are not able to. Figure 4.23 shows a scan where we drove both piezos at the highest voltages we can scan (limited by our high-voltage amplifiers (HVAs)). The piezos are scanning in opposite directions for the same polarity, thus doubling the total scan length range.

Our piezos have a free stroke length of $1.5\ \mu\text{m}$ for a voltage from plus to minus 320 V; with two piezos at our maximum scan in figure 4.23, we should still be able to scan about $0.8\ \mu\text{m}$. Yet, we struggle to scan an FSR, which is $0.39\ \mu\text{m}$. Why are the piezos not scanning as specified? Initially, we presumed that our switch to narrower piezos may have caused this, even though the scan range of a shear piezo does not change as it gets narrower. We refuted this later when we saw very small amounts of thermal-cured epoxy on the sides of both piezos. We suspect that this restrained their movement. Normally, the piezos would be totally unable to move if they were glued to the mount. Nevertheless, if a small amount of epoxy partially crept up the sides as the curing epoxy heated up, then this small amount may be enough to stop the lower regions of the piezo from moving but not the top regions. A future assembly should therefore allow for the possibility of thermal epoxy creep at the piezos.

Epoxy creep may also explain the large non-linear piezo response to the driving voltage that figure 4.23 shows. If we use the first peak from the left to define a frequency scale for the whole graph, then $\Delta f_{\text{FSR}} = 1.072\ \text{THz}$, giving $L_{\text{cav}} = 139.8\ \mu\text{m}$. If, instead, we use the second peak, then we get $2.144\ \text{THz}$ and $69.9\ \mu\text{m}$. Our observed nonlinearity is far larger than that previously observed in the same model of piezo in our research group [55, 125, 145], suggesting the cause is mostly the epoxy on the sides and not just the intrinsic nonlinearity of these piezos. Both values are far away from the microscope measurement of $L_{\text{cav}} = (90 \pm 2)\ \mu\text{m}$. Although, the left and right peaks are equally distant

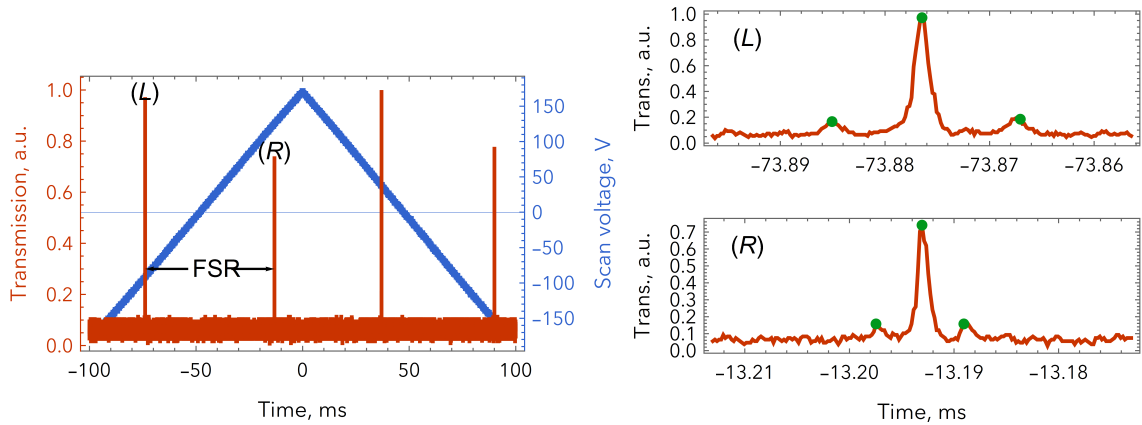


Figure 4.23: Scanning an FSR for the assembled fibre cavity. Here, each peak is the fundamental transverse mode. The two right graphs show the two leftmost peaks in the left graph, and demonstrate a non-linearity of the piezo shearing behaviour: they are at different timescales for the same sideband modulation frequency. The sidebands for both peaks are 270 MHz apart, highlighted with green dots (as well as the carrier peak). We show both the upstroke and downstroke of the scan (blue) to ensure that the observed resonances are not spurious or changing significantly during scans.

from the centre of the scan at 0 V. If we then assume the nonlinearity affects the left- and right-peak frequency scales by the same magnitude, then we can estimate the length as $L_{\text{cav}} = (139.8 + 69.9)/2 = 104.9 \mu\text{m}$. Although still not agreeing with our microscope value to within error, this systematic error is similar in magnitude to the previously observed nonlinearity in our research group. What we find from our estimate, therefore, is that controlling epoxy creep is crucial to controlling the performance of the piezos, just as it is for the cavity itself.

Mechanical stability of the cavity

The second, more crucial, issue with our assembled cavity is that mechanical vibrations impede us from being able to stabilise the cavity length to an atomic transition (‘lock the cavity’) as intended. This is not due to the FFPC itself, but rather due to the mount and vacuum chamber designs. Therefore, we defer discussion of mechanical stability to chapter 5.

Name		Value	Units
Transmission (single mirror)	T	23 ± 1	ppm
Losses (single mirror)	$\mathcal{L}_A + \mathcal{L}_C$	12 ± 1	ppm
Fibre core waist	w_f	3.6	μm
Radius of curvature	ζ	302/345	μm
Feature diameter	D_ζ	33.8/42	μm
Decentration	d	2.5/2.5	μm
Length	L_{cav}	90 ± 2	μm
Mode waist	w_0	5.26	μm
Mode waist at mirrors	w_m	5.75/5.60	μm
Mode volume	V_{eff}	$(39 \pm 1) \times 10^2$	μm^3
Fibre-tip diameter		125	μm
Mode-matching	ϵ	0.29 ± 0.02	
Linewidth	$\Delta\omega_{\text{FWHM}}/2\pi$	18.86 ± 0.02	MHz
Free spectral range	$\Delta\omega_{\text{FSR}}/2\pi$	1.67 ± 0.04	THz
Birefringent splitting	$\Delta_B/2\pi$	13.0 ± 0.3	MHz
Finesse	\mathcal{F}	88 ± 3	$\times 10^3$
Atom-cavity coupling	$g_0/2\pi$	73 ± 2	MHz
Cavity field decay rate	$\kappa/2\pi$	9.43 ± 0.01	MHz
Atomic amplitude decay rate	$\gamma/2\pi$	3.0325	MHz
Co-operativity	C	94 ± 5	

Table 4.3: Summary of the cavity parameters. Value ‘ x/y ’ indicates values for the left and right fibre-tips respectively, identified as B24 and A8. All values assume the wavelength of light is $\lambda = 780.24$ nm, which is the ^{87}Rb D_2 line and will be the subject of our experiments in the following chapters. The corresponding value of γ is from [88].

Summary of cavity parameters

With the characterisation measurements taken thus far, we now summarise the parameters of our assembled FFPC in table 4.3. We group them into the mirror properties, cavity geometry, cavity spectral values and the relevant parameters for CQED. From the latter, we see that our FFPC meets the specifications for a deterministic quantum interface: it is long enough for optical tweezer compatibility, and is well into the strong coupling regime with its high co-operativity (for the best case scenario of an atom located at an antinode).




Figure 4.24: **FFPC** in vacuum cell. The purple tint on the cell arises from the anti-reflection (AR) coating, which we return to in section 5.2.3.

Photo of cavity in vacuum

To fully summarise the completed work in this chapter, we highlight the end product. Figure 4.24 shows our assembled cavity in the vacuum cell. Later in chapter 5, we will refer to this photo when discussing the vacuum assembly.

4.6 Conclusion

In this chapter, we have successfully assembled an **SM-SM FFPC** and placed it into **UHV**. No other group has assembled an **SM-SM FFPC** in vacuum due to the stringent alignment requirements; in doing so, we have solved important challenges relevant to other microcavity experiments. Namely, we have assembled a cavity that is pre-aligned in vacuum for compactness and stability, and using our carefully devised **FFPC** assembly method, we have ensured it remains aligned to sub-micron precision—an essential quality for similar microcavity setups, and one which will only become more important as cavities shrink further and further in search of stronger coupling of matter to the cavity field. To this respect, we have ensured our cavity has enough optical access for trapped atoms held in optical tweezers. We have mitigated possible misalignment in the epoxy-curing stage of the assembly too,

furthering our cause.

To summarise, the strong atom-cavity coupling properties and optical access of our assembled [FFPC](#) may make it suitable as a deterministic quantum interface, using the techniques described in [chapter 2](#). The immediate next step is to put atoms inside it and demonstrate atom-cavity interactions at the quantum level, which we correspondingly progress onto now in the next chapter.

Purcell-enhanced few-atom fluorescence in a fibre cavity

In the preceding chapters we have devised a fibre cavity that is suitable for a fully deterministic quantum interface. To test it, we now want to put atoms inside it.

In this chapter, we present in full our design for our envisaged quantum interface with single atoms trapped inside the cavity with optical tweezers. We realise a compact design, where we load the tweezers with a source of cold atoms near to the cavity. Our tweezers make use of an [SLM](#) to produce reconfigurable 2D arrays of trapped atoms. Furthermore, we can position and move each atom independently of one another with micron accuracy. So far, no experiment has combined a reconfigurable 2D array of atoms with a cavity for strong atom-light coupling.

To demonstrate the advantages of our new design, we perform atom-photon coupling experiments with our [FFPC](#). Firstly we load atoms into the cavity by letting them fall in under gravity [[146](#)]. Secondly, once loaded, we look for the Purcell enhancement of the atomic fluorescence while laser-driving them, allowing us to quantify how strong the atom-photon interaction with our cavity is. From there, we can quantify the properties of the atom loading, and the resulting atom-cavity interactions. Afterwards, it would become possible to combine the cavity with the tweezers.

In this chapter, we firstly give a quick overview of our experimental setup in section 5.1. The first main section 5.2 describes the overall experiment design, both setting the design of the tweezer experiment and explaining the experiment with atoms falling into the cavity that we perform. We describe the unique stability and stabilisation method for our cavity in section 5.3. Then, section 5.4 introduces our two experiments, which we then present and discuss the results of in sections 5.5 and 5.6. Finally, we conclude in section 5.7.

5.1 Experimental map and overview

Figure 5.1 shows a map of our experiment, featuring all the main components. A high-level overview of it is as follows. We place the FFPC inside an evacuated glass cell. Our design allows a 2D array of optical tweezers to be focused at the cavity mode, where each tweezer contains a single atom loaded from a nearby ^{87}Rb magneto-optical trap (MOT) into the cavity. It allows also for probabilistic loading without the tweezers. We closely tune the cavity's resonance to the ^{87}Rb $5^2\text{S}_{1/2} \longrightarrow 5^2\text{P}_{3/2}$ transition at wavelength $\lambda = 780.24 \text{ nm}$ [88],¹ which in shorthand notation is the D_2 transition.² Laser light tuned to the same D_2 transition strikes through the open side of the cavity, interacting with the atoms. Thus, the atoms scatter light into the cavity mode. Cavity light then couples into the output fibre, and we detect that light. The atom's presence in the cavity will cause the detection profile to change: from this change, we can determine how the atom and cavity are interacting.

¹ For quick reference, this is $f_c = \omega_c/2\pi = 384.23 \text{ THz}$.

² It is called D_2 because of the names of the doublet in sodium were, when first discovered by Fraunhofer, D_1 and D_2 . It does not correspond to the term symbol.

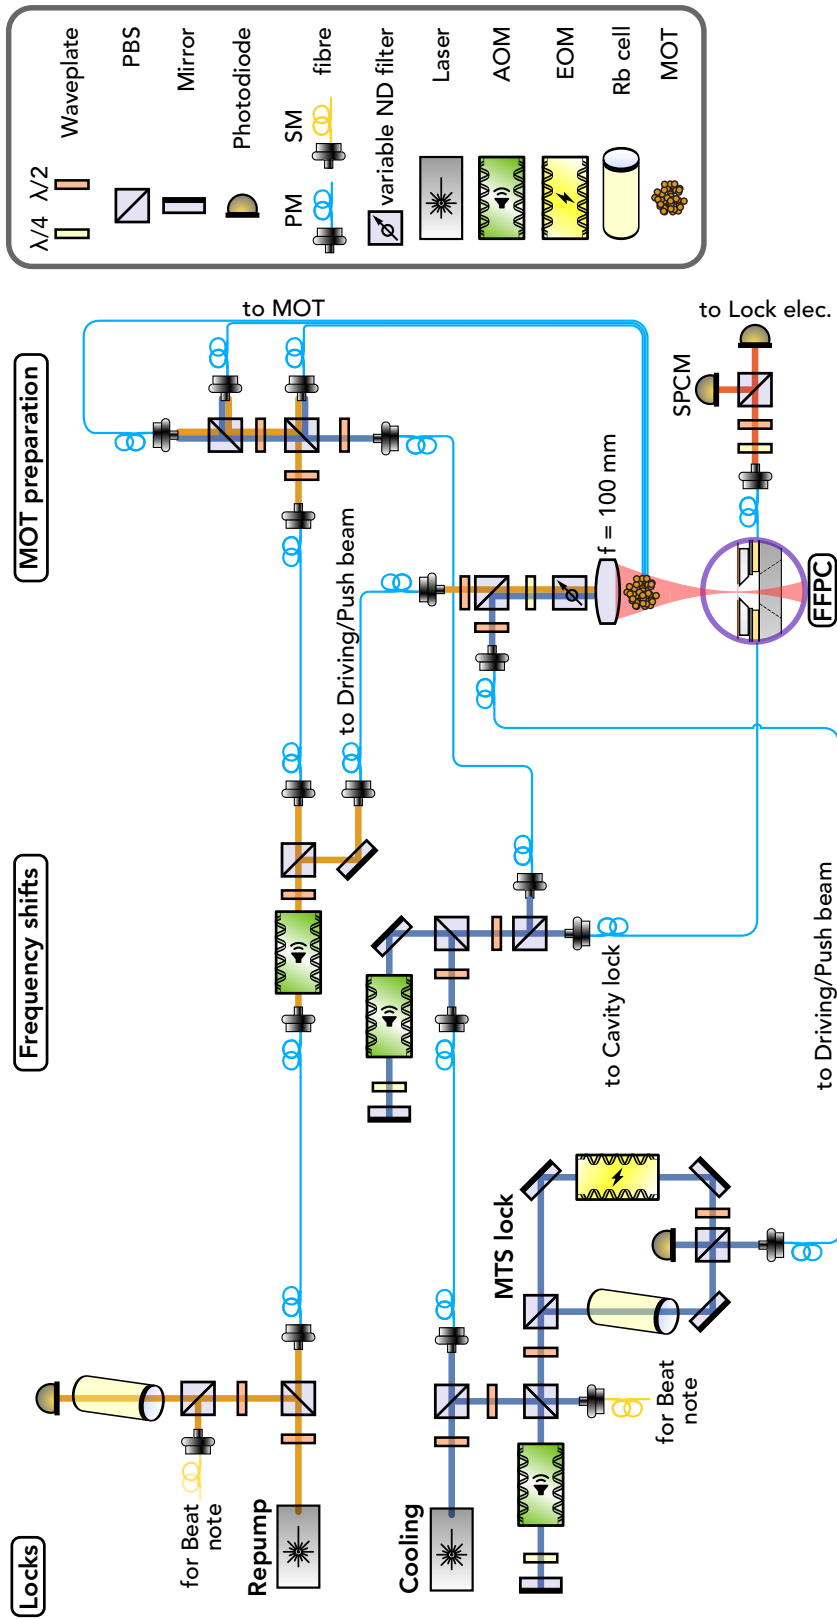


Figure 5.1: Setup used for measuring atom-cavity interactions. We lock the MOT cooling laser frequency to rubidium, and then lock the repump laser frequency to the cooling laser with a beat-note offset lock. Acousto-optic modulators (AOMs) shift the frequencies of the lasers and also enable fast switching on/off. Some cooling and repump light drives the ^{87}Rb D_2 transition within the cavity, scattering light into the cavity mode; in the experiment in section 5.5, we also use this light to push atoms from the MOT towards the cavity. Some cooling light is also used to lock the cavity.

5.2 Experimental design and components

5.2.1 Optical tweezers

Resuming our discussion at the beginning of chapter 4, we intend to trap single atoms at the cavity mode using optical tweezers that are focused onto the cavity mode with a high-NA lens, as we had illustrated in figure 4.1. We know we need a tight focus: we can see precisely how tight by relating NA, which is equivalent to the divergence Θ_{div} ,³ to w_0 [42]:⁴

$$\tan \Theta_{\text{div}} = \frac{w_0}{z_R} \quad (5.1)$$

$$\therefore \text{NA} = \sin \Theta_{\text{div}} = \sin \left(\tan^{-1} \left(\frac{\lambda}{\pi w_0} \right) \right), \quad (5.2)$$

where we have used equation 2.10. This tells us that if we want to position atoms at a cavity standing-wave antinode for the strongest coupling—that is, with at least $\lambda/2$ precision—we need $w_0 > \lambda/2$ and therefore $\text{NA} > 0.54$.⁵ Now we understand why we placed such strict requirements on our FFPC’s optical access.

Because we need to laser-cool the atoms before loading them into the optical tweezers (see appendix C), we load them with cold atoms from a MOT. We shall discuss the principles of the MOT operation in the next section 5.2.2. All we need to know for now is that a cloud of atoms at temperature $T \sim 0.1$ mK are formed at the intersection of 3 orthogonal laser beam paths. Figure 5.2 shows our arrangement of tweezer and MOT beams. The tweezers are focused at the cavity mode with our high-NA lens,⁶ and we simply let the atoms fall by gravity into the cavity to load them. Although two of the beams are not orthogonal to each other, we can consider the vector components of these beams along the Cartesian axes: that

³ See figures 2.3 and 4.1 for relevant illustrations.

⁴ In [42], Meschede uses the small-angle approximation. For these tightly-focused beams, we cannot apply that.

⁵ This also defines the minimum tweezer separation distance.

⁶ Special Optics 54-36.5-30 @780-1064nm (custom design)




Figure 5.2: Laser beam arrangement for the cavity in the vacuum cell. We focus the tweezer beam (yellow) at the cavity mode using high-NA lens(es) close to the cell. The MOT beams (red) intersect 1 mm directly above the cavity mode. If we turn off the MOT, cold atoms will fall by gravity into the cavity. Note that the beams must avoid the corners of the cell, whereupon passing through them may cause undesirable reflections.

is, the axes that are orthogonal. With a more acute angle between the beams, we can form a MOT closer to the cavity for better loading.

We focus on the cavity implementation in this thesis; details on the trapping theory and hologram implementation for the 2D array are available in appendix C and in the works of our research group [107, 108]. Among that work is also a measurement and discussion of the trap lifetime, which is limited by radiative scattering by the trapped atom. The measured trap lifetime was 8.2 s when $P = 96.8$ mW [49]. This is orders of magnitude longer than the decay lifetime of our FFPC, $\tau_c = 1/2\kappa = 8.2$ ns. As such, we can expect to perform many experiments with the same atom before it is lost from the trap.

We only need one lens for the tweezers. Yet, we may want to future-proof the experiment. For this reason, another identical lens and SLM is available to use on the opposite side of the cell to the first lens, as shown in figure 5.2. For example, we can focus STIRAP pulses on single atoms of our choice, rather than driving all of them at once from the side. We will return to this future-proofing in chapter 6.

Lastly, we discuss how we implement the **FFPC** cavity mount that we showed in figure 4.16 into the tweezers setup in figure 5.2. Our stainless steel cavity mount has a conical aperture that allows the strongly-focused tweezer beam to pass through the mount without heating it up. Although, if the cone angle is too large, the cavity mount would be too thin to manufacture. We therefore strike a balance by making the cone angle just a few degrees larger than we need for a $NA = 0.6$ beam to pass through, while also using V-grooves that are small in height and thin, so that we minimise the distance between the **FFPC** and the conical aperture.

To summarise this section, we have demonstrated our design for trapping reconfigurable 2D arrays of atoms at the cavity mode, for which we use optical tweezers. Our design meets the constraints of our setup (optical access, cavity loading, beam arrangement).

5.2.2 Magneto-optical trap

For our source of cold atoms that we load the tweezers/cavity with, we use a **MOT**, a now-standard tool for cold atom experiments. In a sentence, a **MOT** combines laser cooling with magnetic field trapping to produce a dense cloud of atoms at a very low temperature. Typically, the **MOT** temperature is of millikelvin order. For the laser cooling component, three orthogonal pairs of counter-propagating beams intersect to form what is known as optical molasses.⁷ The atoms greatly slow down due to a frictional force, as if they are moving through treacle, thus cooling them. For the trapping, we use a quadrupole field to push atoms towards the centre of the molasses. In this section, we firstly explain the physics of these two components of a **MOT**, so that we can better understand the design constraints of our experiment in the latter part of this section.

⁷ American English for treacle. Similarly, molasses is a singular word, not a plural.

Optical molasses

Optical molasses works by Doppler cooling. Figure 5.3a illustrates the Doppler cooling mechanism in the 1D case. A two-level atom with transition frequency ω_a moves with velocity v_x towards a laser beam of frequency ω_L propagating in the opposite direction. Suppose the laser is red-detuned by Δ_L from ω_a . Then, in the reference frame of the atom, the atom observes a Doppler-shifted frequency of the laser [147],⁸

$$\begin{aligned}\omega_{L, \text{obs}} &= \omega_L \left(1 + \frac{v_x}{c}\right) = (\omega_a - \Delta_L) \left(1 + \frac{v_x}{c}\right) \\ &= \omega_a + \frac{\omega_a v_x}{c} - \Delta_L - \frac{\Delta_L v_x}{c}.\end{aligned}\quad (5.3)$$

Let us assume $\Delta_L \ll \omega_a$. Then, the last term is negligible. So, if we set

$$\Delta_L = \omega_a \frac{v_x}{c}, \quad (5.4)$$

we get

$$\omega_{L, \text{obs}} = \omega_a. \quad (5.5)$$

Therefore, the laser will only resonantly drive atoms moving with v_x .

When the atom absorbs a photon, it will recoil slightly in the $-x$ direction because of the photon's momentum. According to conservation of momentum, the recoil is

$$\Delta p_x = -\frac{h}{\lambda}, \quad (5.6)$$

where h is Planck's constant. Later, the excited atom spontaneously emits a photon after a time $\tau = 1/\Gamma$, where $\Gamma = 2\gamma$ is the atomic linewidth, with an associated recoil. But, after

⁸ We work in the non-relativistic limit, where $v_x \ll c$.




Figure 5.3: Laser cooling of an atom. **(a)** Doppler cooling mechanism. As the atom moves towards the laser, it observes a higher laser frequency. If in turn the laser is red-detuned to compensate for this Doppler shift, the atom sees a resonant laser. The recoil from absorbing a laser photon slows the atom down, cooling it. **(b)** Optical molasses. If we apply six such beams to the atom, one for each direction, then the atom is laser cooled no matter what direction its initial velocity is in.

many absorption/spontaneous emission cycles, the average recoil due to only the spontaneous emission is zero. This is because spontaneous emission is isotropic. The absorption from the laser, however, is not. According to this, the atom experiences a force corresponding to the number of absorption recoils per cycle:

$$F_x = \frac{dp}{dt} = \frac{\Delta p_x}{\tau} = -\frac{h\Gamma}{\lambda}. \quad (5.7)$$

So far, we have neglected the effect of stimulated emission on this process. Let us consider in addition that the intensity of the beam is $I > I_{\text{sat}}$, where I_{sat} is the saturation intensity. Then, for this saturating beam, the absorption and stimulated emission rates are both much greater than the spontaneous emission rate [148]. Since the absorption and stimulated emission rates become equal, the atomic population reaches a steady state with half of the population in the excited state and half in the ground state.⁹ An absorption/stimulated emission cycle yields no net momentum change because the recoil from absorption and from the stimulated emission in the direction of the laser cancel each other out. Therefore, as the

⁹ This is valid both with the (classical) Einstein treatment and with the (quantum) Rabi treatment. For the latter, the intense resonant laser means we oscillate quickly between $|e\rangle$ and $|g\rangle$ (cf. equation 2.38), so the steady-state population in $|e\rangle$ is simply the average of \cos^2 over many cycles, which is $1/2$.

atom is in the excited state only half as often, an absorption/spontaneous emission cycle only happens half as often and so we simply halve the resulting force:

$$\implies F_x = \frac{dp}{dt} = \frac{1}{2} \frac{\Delta p_x}{\tau} = -\frac{h\Gamma}{2\lambda}. \quad (5.8)$$

Note that we must keep track of the factor of one half throughout the following calculations.

How long does it take for the atom to slow down to its minimum speed? First, we calculate the number of cycles it takes to do so:

$$N_{\text{stop}} = \frac{1}{2} \left| \frac{mv_x}{\Delta p_x} \right| = \frac{mv_x \lambda}{h}, \quad (5.9)$$

where m is the mass of the atom, and then we simply multiply by the emission time, giving

$$t_{\text{min}} = N_{\text{stop}} \tau = \frac{mv_x \lambda \tau}{2h}. \quad (5.10)$$

For a gas, we can relate the temperature to an atom's velocity using the standard relation (in the 1D case) [149]

$$\frac{1}{2} mv_x^2 = \frac{1}{2} k_B T \quad (5.11)$$

$$\therefore v_x = \sqrt{\frac{k_B T}{m}}. \quad (5.12)$$

So for one ^{87}Rb atom at room temperature, $v_x \sim 100 \text{ m s}^{-1}$. Thus, inputting this into equation 5.10 yields a Doppler cooling time $t_{\text{min}} \sim 1 \text{ ms}$. This gives us a good idea of how long the cooling process takes.

At this point, we must consider that the atoms we wish to cool will not all have the same velocity at temperature T . Thus, we would be cooling a proportion of atoms in a specific

range of velocities. Now, if we were to set Δ_L according to equations 5.4 and 5.12, we would have $\Delta_L = 35\Gamma$. This would be problematic, because a further analysis shows that $F_x \propto 1/\Delta_L^2$ [150]. Too far detuned, and the cooling force may be too weak to allow us to trap the atoms. Too close to resonance, and the large proportion of atoms that do not see a Doppler-resonant beam will not be cooled, nor would we be able to cool atoms quickly enough according to equation 5.10. Therefore to optimise this interplay, Δ_L was previously found empirically by Holland for similar parameters as we will use [49]. We use $\Delta_L = 2.25\Gamma$, which is also $\Delta_L/2\pi = 13.5$ MHz (using 2γ from table 4.3).

Of course, the atom cannot slow down to zero velocity by Doppler cooling. This is because Doppler cooling only works if $\omega_{L, \text{obs}} = \omega_a$. Yet, due to the uncertainty principle, ω_a has an uncertainty $\Delta\omega_a$ due to its natural linewidth Γ . Therefore, by Doppler cooling, we cannot reduce the energy of the atom below

$$E_{\min} = \frac{1}{2}\hbar\Delta\omega_a = \frac{1}{2}\hbar\Gamma, \quad (5.13)$$

which we divide by k_B to yield the Doppler-limited temperature,

$$T_D = \frac{\hbar\Gamma}{2k_B}. \quad (5.14)$$

So, we can expect the atom, and therefore the MOT, to have a temperature of 146 μK [88].¹⁰

Optical molasses is the 3D version of this Doppler cooling arrangement. Going from 1D to 3D, our analysis above is identical for the y and z axes. We therefore cool atoms by using six beams: one for each direction, as shown in figure 5.3b. At their intersection point, no matter what direction the atom's velocity is in, it will always be Doppler-cooled. In the laboratory,

¹⁰ Although we have calculated this in the 1D case, we obtain the same value of T_D in 3D, because firstly we have $v = \sqrt{v_x^2 + v_y^2 + v_z^2} = \sqrt{3}v_x$, and secondly the right hand side of equation 5.11 is multiplied by three as there are three degrees of freedom. The resulting two factors of $\sqrt{3}$ cancel each other out.

we construct optical molasses by overlapping three beams and then retroreflecting each beam to get the remaining three counter-propagating beams.

In addition to the Doppler cooling mechanism we have discussed, there are also sub-Doppler cooling processes [150]. In particular, we experience polarisation gradient cooling in our MOT [151]. For brevity, we omit discussion of this effect, but the net outcome is that our MOT temperature is $T \approx 100 \mu\text{K}$, which is below the Doppler limit.

Quadrupole magnetic field

Optical molasses cools the atoms, but does not trap them. To load the atoms into the cavity, we must first localise them if we want to load them efficiently. Therefore, we create a quadrupole magnetic field by adding a pair of current-carrying coils, as shown in figure 5.4a. These coils, which we hereafter refer to as MOT coils, are in an anti-Helmholtz configuration, where the coil currents run in opposing directions. For this configuration, the zero point of the magnetic field is at the centre of the molasses, halfway between the coils. So, the further away an atom is from the centre, the stronger the field that it feels—similar to the gravitational field from a point mass.

How does this magnetic field trap the atoms? By the Zeeman effect, it lifts the degeneracy of the M_F sublevels of the hyperfine energy levels of the atom. Then, because this splitting is position-dependent, we take advantage of it to construct a restoring force towards the field's zero point. Here, a hyperfine level is denoted F , and is a quantum number for the total angular momentum of the atom, where its z component¹¹ is M_F .

Figure 5.4b shows how Zeeman splitting changes the hyperfine level structure of the atom. For the right atomic structure, as we have for ^{87}Rb , we have polarisation-dependent transitions to different Zeeman sublevels. Therefore, the shift in energy levels gives rise to a conservative

¹¹ For atomic physics, we can usually arbitrarily choose the axes such that we simplify our problem. Here, we choose the z axis to coincide with the direction of the applied magnetic field.

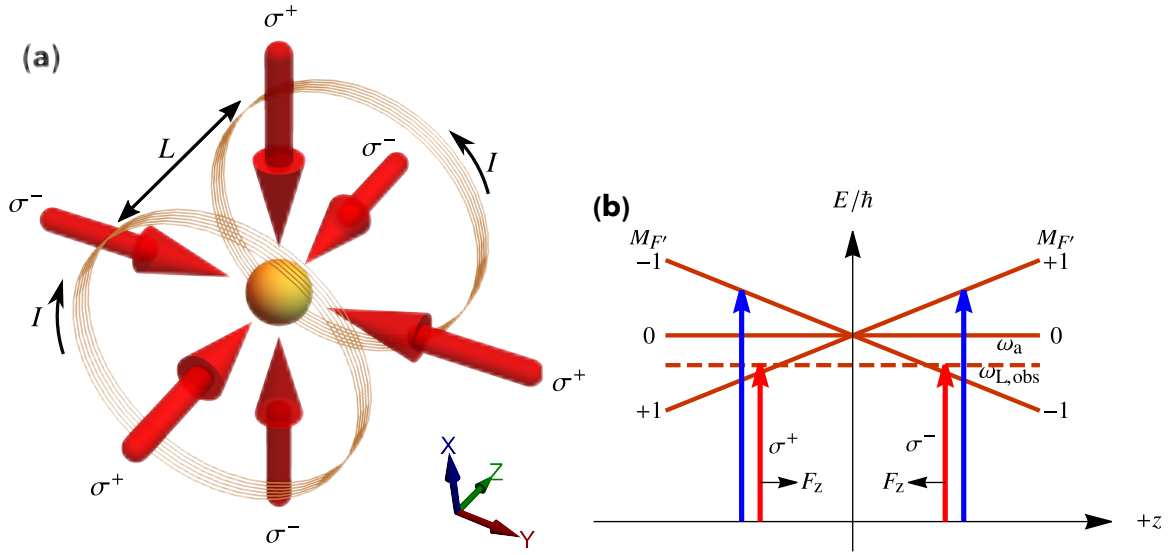


Figure 5.4: Using a quadrupole magnetic field to trap atoms. **(a)** A MOT. From optical molasses (figure 5.3b), we add two current-carrying coils. By driving them with oppositely directed currents (anti-Helmholtz configuration), we create a quadrupole magnetic field. The field is zero where the atom is in this picture (halfway between the coils), and increases in strength at greater distances from this point. The effect of the field is to push the atom towards the zero point. σ^+ (σ^-) denotes right(left)-handed circular polarisation respectively. **(b)** Zeeman splitting by the field. If we set the beam polarisations accordingly (red arrows), then an atom at $+z$ ($-z$) will interact more strongly with the $M_{F'} = -1$ ($+1$) sublevel. The levels higher in energy become further off resonance and are thus driven more weakly (blue arrows). Altogether, this results in a net recoil F_z that pushes atoms towards $z = 0$.

potential,¹² with a corresponding trapping force

$$F_z = -\alpha_{\text{MOT}}|z|, \quad (5.15)$$

where we implicitly choose $z = 0$ at the field's zero point, and α_{MOT} is a constant incorporating the magnetic field strength gradient,¹³ atom, laser parameters etc [152].

To see this restoring force, first we must set the beam polarisations. We see from equation 5.15 and figure 5.4b that we should set them as follows: σ^+ for the beam propagating

¹² Just like with the tweezers: see appendix C.

¹³ For an anti-Helmholtz configuration, the field strength gradient along the radial directions (here, x and y) is double that along the axial direction (here, z).

towards $+z$, and σ^- for the counter-propagating beam towards $-z$, where σ^+ (σ^-) denotes right(left)-handed circular polarisation respectively.¹⁴ Because the beams are red-detuned, an atom at $-z(+z)$ will interact more with the $M_F = +1(-1)$ sublevel. Thereby, there is a resulting imbalance in recoils, and the net recoil always pushes the atom towards $z = 0$. The same argument applies for the x and y directions.

We must make one remaining correction to our MOT configuration. Off-resonant excitation followed by spontaneous emission can leave the atom in a different ground state than the one we intend. For our level scheme, this occurs roughly once in every 100 cycles [153]. If indeed this occurs, then atoms in this state can no longer be excited with the laser. Hence, we call this a ‘dark’ state. Atoms remain ‘dark’ unless laser-driven back to the excited state at the corresponding frequency. Thereafter, the cooling transition can keep cycling again. So, as well as the cooling laser, we require an additional laser for the dark atoms, which we call the repump laser. In the MOT, we just input the repump beams at exactly the same parts we input the cooling beams, as shown in the upper-right of figure 5.1.

Referring to figure 5.2, we now understand that we can reasonably expect to form a cloud of cool atoms at $T \sim 100 \mu\text{K}$ at 1.0 mm above the cavity. We also see that a stronger field gradient (which corresponds to larger α) can give us both a larger capture range (of velocities) and a more compact MOT. These aspects are crucial for our design where we need to efficiently load atoms into the cavity, and also where the MOT parts must fit around the cell/cavity. With respect to this, we now discuss the design of the magnetic field coils and the molasses beams.

¹⁴ Precisely, σ^+ denotes atomic transitions where $\Delta M_F = +1$, and σ^- where $\Delta M_F = -1$. For our chosen quantisation axis and propagation direction, these are just right(left)-handed circular polarisations.




Figure 5.5: Photo of the MOT apparatus. Here, we see the diagonal MOT beams (see figure 5.2); inside the tube and on the retroreflector are the necessary optics for the MOT. The blue cube wireframe contains the six compensation coils that allow us to precisely position the MOT. The tweezer lens(es), not pictured here, can occupy the space above and below the cell.

Design of coils and beams

Firstly, we show the MOT apparatus with the coils and beams in figure 5.5. Then, we discuss the design of each in turn.

Let us start firstly with the design of the MOT coils. We desire three properties: a sufficiently strong field gradient, compactness, and the ability to run cool. There is an interplay between these properties: if we need a stronger field, we could increase the current or the number of turns, but then the coils may overheat, causing the cavity length to drift.¹⁵ Similarly, the coils should be brought close together for a strong gradient, but not too close to or touching the vacuum chamber, wherein the cavity would also heat up. We do not need to be particularly precise with this interplay: just within the right parameter range, which we now set.

¹⁵ More turns also increases the time for the coils to switch on and off, but usually we can account for this by making our experimental sequence a few milliseconds longer: see section 5.6.3.

Firstly, we set a reasonable field gradient to aim for. Based on previous working experiments with MOTs in our research group, we aim for a field gradient $\gtrsim 1 \text{ mT cm}^{-1}$ since we are making a similar sized MOT [55, 153]. Secondly, as the cell diameter is 27 mm, we can set an inner coil diameter $2r_{\text{in}} = 60 \text{ mm}$, giving us leeway around the cell when aligning them. Thirdly, to improve on previous experiments, we use a similar number of coil turns to our group's previous water-cooled coils, but operating at a much lower current so that we can passively cool them, thereby keeping the coils compact. We put these three points together to specify our coils as follows.

Close to $z = 0$, the field gradient for these coils is constant, and is [154]¹⁶

$$\left. \frac{dB}{dz} \right|_{z=0} = \mu_0 N I \frac{3r^2(L/2)}{\left(r^2 + (L/2)^2\right)^{\frac{5}{2}}}, \quad (5.16)$$

where μ_0 is the magnetic constant, N is the number of coil turns, I is the current, r is the radius of the coils and L is the separation distance between them. If we choose $N = 159$, $r = 71.5 \text{ mm}$, $L = 48 \text{ mm}$ and $I = 3.3 \text{ A}$, we get a gradient of 0.99 mT cm^{-1} , as intended. Our coils use 18AWG wire.¹⁷

As well as the MOT coils, in figure 5.5 we also see a cube wireframe. These hold the compensation coils, whose purpose is to compensate for residual magnetic fields (e. g. from the Earth) at the MOT formation point and also to allow us a few millimetres of freedom in positioning it exactly where the molasses beams overlap. Three pairs of square coils allow us to compensate along x , y and z . Because the compensation fields are much weaker, our only design constraint is that they must fit around the experimental apparatus. We hand-coil them,

¹⁶ This is derived by starting from the Biot-Savart law for a current-carrying loop, calculating the field from both coils, and then Taylor expanding to second order. To suppress second-order terms in the gradient and therefore keep the field linear in the vicinity of the MOT (see figure 5.4), we use $L = \sqrt{3}r$ —although a linear field is not a strict requirement for creating a MOT.

¹⁷ American Wire Gauge, about 1.02 mm.

using 9 turns.

Our design allows to quickly and easily 3D print the coil holders ('formers'), and to slide the MOT apparatus in and out of the setup on an optics breadboard/rails so that we may rapidly modify it. Unfortunately, although we had MOT coil holders ready to print, we were not able to coil them due to time constraints of the workshop performing the coiling. Therefore, we used our printed compensation coils along with a set of MOT coils from the previous experiment of Jerome Dilley [155]. These have, for $I = 6.2$ A, a gradient of 0.49 mT cm⁻¹. Fortunately, we were still able to form a MOT without active cooling of the coils, and did not notice any coil-caused temperature drifts of the cavity at this current. Although, we then find the MOT is slightly too large for optimum cavity loading, as we will discuss in section 5.6.1. In future experiments, we may well want to update the design of the MOT, as we will briefly discuss in chapter 6.

Now, let us discuss the design of the molasses beam setup. We desire three properties: equal powers of all six beams ('balancing') so that there is no net force away from the trap, compactness, and good control over the beam size. We get all three properties at once by injecting the molasses light through an assembly of tubes. We input linearly polarised light from PM fibres into three tubes. Figure 5.5 shows the tubes; the top-right of figure 5.1 shows how we prepare linearly polarised light to send into the tubes. Inside each tube is a lens, a quarter-wave plate (QWP) and an aperture, giving us collimated, circularly polarised beams with well-controlled diameters. We attach QWPs onto the retroreflectors to give the correct polarisations for the counter-propagating beams. We mount both the tubes and the retroreflecting mirrors on thick steel posts well away from the cavity, giving good mechanical stability. Thus, we satisfy all three properties. One further improvement can be made with regards to balancing the beams: when striking through the AR-coated cell at an angle, still about 2% of incident light is reflected. Although the beams are still well-balanced enough for

us to form the molasses, we have extra space in the tube reserved for a polarising beam-splitter (PBS) and half-wave plate (HWP), should we need to split the beams and balance them even better in the future.¹⁸

With respect to the loading of the tweezers, an atom falling 1.0 mm under gravity will increase its thermal velocity by ≈ 0.07 mK. Therefore, given our trap depth $T_d = 47.4$ mK from appendix C, an atom falling from the MOT is likely to be caught by the tweezers.¹⁹

5.2.3 Vacuum chamber

With laser cooling, trapping and our MOT described, we now need to provide the rubidium atoms for this experiment. The standard way to do this is to dispense rubidium in a glass cell, as one may use for spectroscopy. However, at normal atmospheric pressure, there are too many energetic atoms that can knock our atoms out of the tweezers and MOT. Cooling all the atoms is impractical, so instead we operate at very low pressure—UHV—and seal off the system from air. Therefore, we need a vacuum chamber. In this section, we firstly show the overall apparatus in figure 5.6, then discuss the characteristics of our vacuum cell and chamber relevant to our experiment. The key quantity is our final vacuum pressure, which is 2.5×10^{-9} mbar for our experiments.

Vacuum cell

The centrepiece of the vacuum chamber is the vacuum cell, since this is where we perform experiments on the FFPC. For a close-up picture of the cell, we refer to figure 4.24. To maintain good optical access, we use a custom-made glass cell²⁰ that makes our arrangement

¹⁸ An alternative approach is to focus the beam slightly using a weakly curved mirror such that the intensity of the counter-propagating fields are equal at the point where the MOT is formed.

¹⁹ This assumes that the atom follows a ballistic-only trajectory, then equates its kinetic energy to thermal energy, $\frac{1}{2}m\mathbf{v}^2 = \frac{3}{2}k_B T$.

²⁰ ColdQuanta CQOX0001D




Figure 5.6: Photo of the vacuum chamber. We describe each component in the main text. Obscured is the valve for the turbo pump; not present are the protective components for the fibres on the air side.

of beams possible. It is compact; its inner dimensions²¹ are $(60 \times 21 \times 21)$ mm, with BK7 AR-coated glass walls 3.5 mm thick. We note that due to refraction by the glass, the diagonal MOT beams and the tweezers are shifted by a small amount,²² so we compensate for this when we align the beams.

The cells walls are (broadband) AR-coated for both 780 nm and 1064 nm light, the two wavelengths we use. However, the AR coating is less effective when the light strikes at an angle: as we had mentioned, we lose about 2% of the diagonal MOT light when entering one side and exiting the opposite side of the cell.

Vacuum assembly for fibres

On discussing our fibre implementation in vacuum, we may have to correct a few assumptions. From chapter 4, one may have assumed the fibre simply consisted of the glass core and

²¹ Of the square part of the cell.

²² MOT beams shifted 0.7 mm closer to the centre of the cell, and tweezer beam extremities shifted by 0.5 mm closer to the centre of the cell.

cladding. We also had not discussed how we inject light into these bare fibres. Now, we address those two points.

Firstly, the fragile glass fibres need to be protected while we handle them, but the typical plastic shroud is not suitable for vacuum due to outgassing. Even when the fibre-tips were originally coated, they needed to be placed in a vacuum chamber. Therefore, the fibres are coated with copper by the manufacturer using a ‘CuBall’ process, similar to electroplating.²³ This leaves a 20 μm -thick copper layer. Not only does this make the fibres UHV compatible, the copper also protects the fibres from breakage (compared to bare glass). However, as the coating is thin, they are still fragile. Moreover, the copper exhibits a shape memory, in which any curves in the fibre are kept under tension. If due care is not taken when the fibre is repositioned, then the built-up tension is released and the fibre can snap. So we still must be careful when handling the fibres.

Secondly, we add connectors to the fibres using a fibre polishing and ‘connectorisation’ kit²⁴. With this kit plus some copper etching solution, we cleave the fibre at the non-coated end, glue a standard FC/APC fibre connector onto it, and then polish the fibre down so that the connector and fibre surfaces are flush. Thus, we can directly connect the fibre to our fibre couplers. When we put the fibres in vacuum, we first cut the connectors off, then install the fibres, and finally add connectors to them again on the air side. There is a limited number of times we can do this, because the fibres are not particularly long: about a couple of metres, since they had to fit inside a small drum for the IBS machine when they were mirror coated. In future, we could splice the fibre. Here, splicing means to join two fibres together by fusing them with high heat from an electric arc.²⁵ Splicing is both less lossy and allows us to make

²³ See [139, appendix A].

²⁴ Thorlabs CK05

²⁵ In particular, this is called fusion splicing; mechanical splicing, where we simply bring the two fibres as close as we can and fill in the gap with a refractive index-matching gel, is both incompatible with vacuum and much lossier.

the fibre as long as we need. On the other hand, splicers are not compatible with metal, so the fibre must be etched to bare glass at the splice point. For our relatively short fibres, the splice point may be in vacuum: overcoming this fragility in vacuum would therefore require us to design a [UHV](#) suitable splice protector. For now, we proceed with the connectorisation kit, where we did not find any significant losses in our final polished connectors.

To install the fibres in [UHV](#), we need to form a vacuum-tight seal around them where they enter. For this, we build our own fibre feedthroughs by combining Teflon ferrules with Swageloks [[156](#)], which are shown in figure [5.6](#). Firstly, we drill a hole (280 μm diameter)^{[26](#)} into a piece of Teflon which fits into a 1/8" Swagelok. Next, we insert the fibre (165 μm diameter) into the Teflon ferrule. Tightening the Swagelok then compresses the Teflon around the fibre, forming a hermetic seal. The main advantage of this approach is the ability to replace the fibres without damaging them, simply by loosening the Swagelok and removing the ferrule. The ferrules themselves are also reusable. This is why we choose this repeatable approach instead of gluing the fibre into the ferrule [[157](#), [158](#)]; we needed to be able to remove the cavity without having to break the fibres, for example if we had to make changes to the vacuum apparatus. We did, however, experience a small disadvantage in terms of reliability: on one or two occasions, we experienced a vacuum leak when moving a fibre, e. g. to connect it to another fibre coupler. The fibre likely slipped slightly. After we re-tightened the Swagelok, the leak stopped. Having used both these reusable ferrules and also the glue-based ferrules before, we believe the reusable ferrules are better suited to [FFPC](#) experiments: their advantages greatly outweigh the small reliability disadvantage. Furthermore, we expect that by optimising the ratio of the hole and fibre diameters, and the amount of torque applied to tighten the Swagelok, we can greatly improve the reliability of our feedthroughs.^{[27](#)} As an aside, we note that all these designs work with multiple fibres in

²⁶ We use the closest imperial equivalent: an 11 Thou drill bit.

²⁷ We used one full turn past finger-tight, as recommended in Abraham and Cornell [[156](#)]. Using aluminium

each Swagelok (one hole for each fibre) [156, 159].

If a fibre breaks too close to the fibre feedthrough, the FFPC may become unusable as the fibre would then be too short to add connectors or splice. Alleviating this possibility, we protect the fibre with three components: a plastic protective box enclosing the fibre feedthroughs, heat-shrink enclosing the fibre, and strain relief posts supporting the fibre along its length.

Rest of vacuum chamber

We attach the FFPC mount to the vacuum chamber using a custom-built holding finger, illustrated in figure 5.7, which also shows how it attaches to the rest of the chamber. The holding finger attaches to a CF40 double-sided flange; the flange attaches to the cell from the front and to a five-way cross from the back. By doing this, we have full rotation of both the cell and the finger independently of one another.

The finger bolt holes are threaded, and so the double-sided flange bolt holes are not. This means that the bolt heads sit on the five-way cross side. Therefore, if we want to remove the finger, we must remove the double-sided flange too. Ideally, we would have the finger not threaded and the double-sided flange threaded, but in any case we did not encounter much difficulty with having to remove both parts together.

In our setup, we did not need full rotation of the finger, so we instead used a fixed flange on the five-way cross. This ensures that the cavity finger is always mounted parallel to the optical table (there are three more possible orientations; in total, four at 90° from each other). We had also used a rotatable flange on the five-way cross for this setup without any issues.

Similarly to the fibres, for electrical connections we use an electrical feedthrough.²⁸

instead of Teflon [159] may possibly also offer better grip on our copper-coated fibres.

²⁸ Lewvac ZFL-LF2S1/8-40CF




Figure 5.7: Holding assembly for cavity mount. **Left:** Cavity mounting finger for vacuum chamber. **Right:** Exploded assembly showing how the finger attaches to the rest of the chamber. Starting from the left, we have the cell, cavity finger, customised double-sided flange, and five-way cross (which is connected to the rest of the vacuum chamber). Both the vacuum cell and the five-way cross can have rotatable flanges, allowing the finger and cell to rotate independently of each other.

Connected to it are the piezo wires, and also two rubidium dispensers²⁹ which slowly sublimate rubidium into the chamber when heated.³⁰ We use one dispenser, and keep the other as a fail-safe. The dispensers and two piezos all share a common electrical ground.³¹ Owing to space restrictions, we connect the wires on the vacuum side using push-fit connectors.³²

Assembly procedure

Our assembly procedure for the chamber is fairly standard [160, appendix B], with some differences for our setup. We placed the cavity finger and double-sided flange in a mounting bracket³³ that is held in a boss clamp, allowing us to hoist the finger up to the chamber and

²⁹ SAES Getters RB/NF/4.8/17 FT10+10

³⁰ They use resistive heating: the higher the electric current, the greater the dispensing rate.

³¹ The two piezos' bases are electrically connected via the steel cavity mount, and hence also the dispensers.

³² Lewvac FHP-BECU-3IL-CON. Barrel connectors from Lewvac were also used but, as well as being large and difficult to fit, also came loose during operation: we observed intermittent jitter in the piezo position before we re-tightened them. Push-fit connectors, however, risk snapping the wire upon connection. We therefore recommend custom-making connectors to fit properly.

³³ Thorlabs VFA275/M

position it precisely. To stop the gasket for the vacuum seal from falling and snapping the fibres during installation, we hooked it with Kapton wire onto the back of the double-sided flange prior. We also positioned the electrical feedthrough using the same mounting bracket and boss clamp. Lastly, we curved/kinked the fibres/piezo wires respectively such that they do not obstruct the MOT beams coming through the viewport.

To reach a low vacuum pressure, we must heat the assembly to a high temperature to release outgassing vapours from the components and chamber walls ('bake'). Because of time constraints, we baked the system without the glass cell or cavity (that is, a blank flange in place of the cell) at 200 °C.³⁴ After pumping down with a turbomolecular pump and ion pump combination,³⁵ we achieved a pressure of 1.0×10^{-9} mbar. After we installed the glass cell and cavity, the final pressure was 2.5×10^{-9} mbar. This is fine for our experiments where we need to make a MOT. If, for example, we need better tweezer performance, we can expect to reach 5×10^{-11} mbar simply by baking the cell and baking for longer, as has been achieved by Holland for a similar setup [49, appendix A].³⁶

5.2.4 Lasers and locking

We use two lasers for our experiment: a cooling laser, and a repump laser. Both are external cavity diode laser (ECDL) models,³⁷ where we obtain both a narrow linewidth of kilohertz order and tunability over several gigahertz by enclosing the laser diode cavity in another larger cavity. By changing the larger cavity's length, we can therefore 'pick out' the desired frequency of the laser cavity.

In this section, we firstly describe the atomic transitions that we use. Secondly, we explain

³⁴ We replaced the Teflon ferrules with stainless steel plugs during this bake.

³⁵ Pfeiffer Vacuum TSU071E; Gamma Vacuum 25S-CV

³⁶ Although untested at 200 °C, we expect the cavity to survive a bake temperature this high without misalignment, according to our results in section 4.4.4.

³⁷ Toptica DL100

how we lock the lasers to these transitions, effectively eliminating unwanted fluctuations in the laser frequency.

Atomic transitions used

Our experiments take place on the $^{87}\text{Rb } 5^2\text{S}_{1/2} \longrightarrow 5^2\text{P}_{3/2}$ transition line, or D_2 line. Its hyperfine level structure is displayed in figure 5.8. We use the D_2 line because it contains a cycling transition for which, when laser-driven, the atom always falls back into the same state after each photon emission.³⁸ Thus, the atom continually emits photons until an off-resonant drive and decay occurs. Here, our cycling transition is $|F=2\rangle \longrightarrow |F'=3\rangle$: by the selection rules for electric dipole transitions [147], an excited atom in $|F'=3\rangle$ can only ever decay to $|F=2\rangle$. We use this cycling transition as our cooling transition. As we discussed in section 5.2.2, we also need our repump laser to keep the cooling transition cycling. For our level scheme, the atom falls into the dark state whenever it is off-resonantly excited from $|F=2\rangle \longrightarrow |F'=2\rangle$ and then spontaneously decays to $|F=1\rangle$. Therefore, our repump laser drives the atom out of the dark state on the transition $|F=1\rangle \longrightarrow |F'=2\rangle$; then once it spontaneously decays to $|F=2\rangle$, the cooling process can continue.

Frequency locking

Consider the frequency locking of a laser to the cooling transition. If unlocked, fluctuations of the laser frequency due to the laboratory environment—e. g. heating, vibrations—will mean we cannot always drive the transition resonantly (or at our chosen detuning). So, to lock it, we must compensate for these fluctuations. Let us assume a fluctuation causes the laser to slightly drift off-resonance from the cooling transition: to correct the erroneous laser

³⁸ Also, we use the D_2 line rather than the D_1 line for this reason, which does not have a cycling transition [88]. Moreover, although ^{87}Rb is less abundant naturally than the other stable isotope ^{85}Rb , only the $^{87}\text{Rb } 5^2\text{P}_{3/2}$ fine level on the D_2 line has an $|F'=0\rangle$ excited hyperfine level. This level is not Zeeman-split in a magnetic field, which allows us to realise a true Λ -level system.




Figure 5.8: Hyperfine levels of the $^{87}\text{Rb } D_2$ line. We use two lasers to drive atomic transitions: a cooling laser and a repump laser, as indicated. Transition frequencies are from [88]. Here, all quantities are in non-angular units, so $\Delta_L = 13.5 \text{ MHz}$ as labelled. In all other sections, we use $\Delta_L/2\pi = 13.5 \text{ MHz}$.

frequency, we need an error signal that directly corresponds to the deviation from resonance. Immediately, we then feed back the error signal to our laser frequency to correct the error. As long as we keep doing this, we keep it on resonance.³⁹

How do we generate the error signal? The standard method is to apply FM to the laser, producing laser sidebands that act as phase markers⁴⁰ to which the reference (resonance) has a fixed phase relationship. If the laser is off-resonance, it would be closer in phase to one marker than the other. We can measure this phase difference to get our error signal.

Consequently, we can gain insight on how to measure this error signal as follows. Suppose, at some fixed position, we have a field

$$E(t) = E_0 e^{i\omega t},$$

³⁹ Note that the error does not have to be with respect to a resonance: any well-defined reference value will suffice.

⁴⁰ In section 4.4.1, we also produced laser sidebands in this way, but instead used them as frequency markers.

that we input to an electro-optic modulator (EOM). We drive the EOM at angular frequency ω_m , which corresponds to modulating the refractive index n harmonically,

$$n(t) = n_0 + A \sin(\omega_m t), \quad (5.17)$$

where A is the modulation amplitude. The modulated wave on the output of the EOM is thus

$$E(t) = E_0 e^{i\omega t} e^{iA \sin(\omega_m t)}, \quad (5.18)$$

where we have dropped a constant phase factor. If we choose the value of A such that most of the power is concentrated in the carrier (i. e. central peak) and first order sidebands only, then we can rewrite this as [148]

$$E(t) = E_0 \left(-A_s e^{i(\omega - \omega_m)t} + A_c e^{i\omega t} + A_s e^{i(\omega + \omega_m)t} \right). \quad (5.19)$$

where A_s and A_c are the amplitudes in the sidebands and carrier respectively.⁴¹ Immediately, we see that we have three frequency modes corresponding to the carrier and two sidebands (like in figure 4.11), and that the two sidebands are in antiphase with each other but one sideband is in phase with the carrier.

Now suppose these three modes interact with the reference, e. g. atoms in a rubidium vapour cell. We model this using a transfer function, $F(\omega)$, accounting for what gets transmitted and what gets absorbed at frequency ω .⁴² In this case, we modify equation 5.19 to get

$$E(t) = E_0 \left(-A_s F(\omega - \omega_m) e^{i(\omega - \omega_m)t} + A_c F(\omega) e^{i\omega t} + A_s F(\omega + \omega_m) e^{i(\omega + \omega_m)t} \right). \quad (5.20)$$

⁴¹ Equation 5.19 is derived in [148] using a mathematical identity to rewrite the exponent with the sine as a Bessel function (including first order terms only). We have then absorbed the Bessel coefficients into the A_s and A_c terms.

⁴² Or reflected, as we shall consider with cavities soon.

Here, $F(\omega)$ is a Lorentzian centred on the atomic transition with natural linewidth Γ . Our detector or photodiode measures intensity, so will see $|E(t)|^2$. So, the resulting expression for the detected signal will contain exponents with DC (i. e. not oscillating), ω_m and $2\omega_m$ terms—but we need only focus on the ω_m terms.⁴³ To calculate these terms only, we use equation 5.20 to find

$$|E(t)|_{\omega_m}^2 = 2|E_0|^2 A_s A_c (\text{Re}(\mathbb{F}) \cos(\omega_m t) + \text{Im}(\mathbb{F}) \sin(\omega_m t)), \quad (5.21)$$

where the ω_m subscript denotes that these are the ω_m terms only, and where

$$\mathbb{F} = F(\omega)F^*(\omega + \omega_m) - F^*(\omega)F(\omega - \omega_m). \quad (5.22)$$

Additionally, we can safely assume $\omega_m \ll \omega$, since electronic frequencies (usually of radio frequency (RF) order) are much slower than optical frequencies. As a result, we can input the Taylor expansion

$$F(\omega \pm \omega_m) \approx F(\omega) \pm \omega_m \left. \frac{dF}{d\omega} \right|_{\omega} \quad (5.23)$$

to equation 5.21 to further simplify it [161]:

$$\implies \mathbb{F} \approx \omega_m \frac{d|F|^2}{d\omega} \quad (5.24)$$

$$\implies |E(t)|_{\omega_m}^2 \propto \omega_m \frac{d|F|^2}{d\omega} \cos(\omega_m t) \quad (5.25)$$

where the sine vanishes on the second line because $\text{Im} \left(\frac{d|F|^2}{d\omega} \right) = 0$. We therefore see that from the detected photodiode signal we can extract the phase difference and therefore error

⁴³ The direct current (DC) terms will not vary in time as the phase changes, so they are not important to us. We do not gain any extra information from keeping the $2\omega_m$ terms either, so we remove them with a low-pass filter (we will discuss this soon).

signal: if we phase-shift the detected signal (to get a sine) and then beat it against the **EOM** driving signal (a sine), we obtain a **DC** signal directly corresponding to the phase difference from resonance, because the transfer function is centred on resonance and is symmetric. We can reinforce our picture as follows: on resonance, the two values of $F(\omega \pm \omega_m)$ are equal, and so the sum of the two sideband phases cancel each other out. Off resonance, the two values of $F(\omega \pm \omega_m)$ are different, and we would measure this as a phase change. After doing this, we eventually derive an error signal corresponding to this phase change [161]

$$\epsilon \propto \mathbb{F} \approx \omega_m \frac{d|F|^2}{d\omega}. \quad (5.26)$$

Figure 5.9a shows ϵ as a function of the Lorentzian peak. It is zero at resonance; if the laser frequency is erroneous, then this signal can be fed back to the laser to correct the error.

Cooling laser lock

For the cooling laser, we generate the error signal using modulation transfer spectroscopy (MTS) [162]. As shown in figure 5.9b and in the bottom-left of figure 5.1, a pump beam and a probe beam fluoresce atoms in a rubidium vapour cell. An **EOM** modulates the pump beam to produce sidebands on it. Then, in the vapour cell, four-wave mixing (FWM) [42, 163] transfers the pump sidebands onto the probe. In **FWM**, two of the three pump frequencies and the probe frequency interact to generate a fourth frequency. This happens for each of the three possible combinations of two of the pump frequencies with the probe frequency, therefore adding sidebands to the probe. Since the probe beam has sidebands and is partially absorbed by the rubidium vapour, it has a field precisely of the form of equation 5.20. Because **FWM** is a resonant process, but our pump beam may be erroneously off-resonant, we simply compare the phase of the pump (**EOM**) modulation signal to the phase of the modulated probe (corresponding to the photodiode signal). The further off-resonant the pump is, the larger the

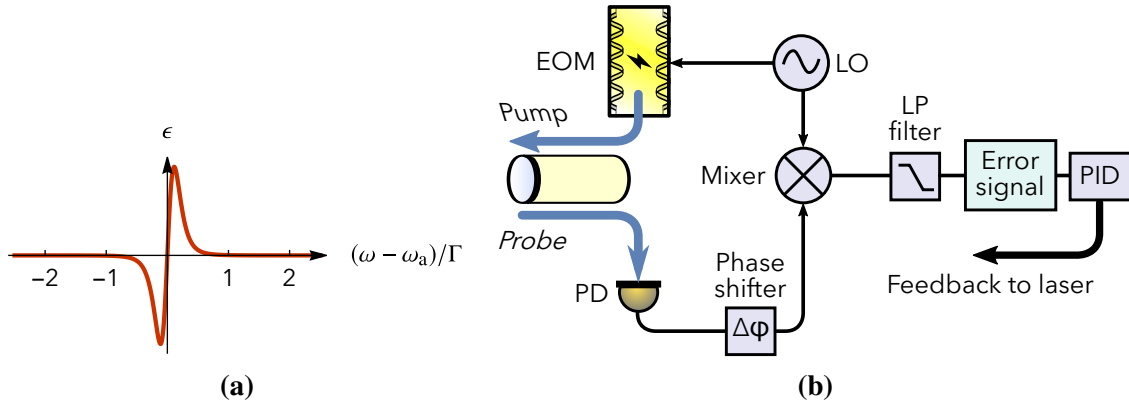


Figure 5.9: Frequency locking using **MTS**. **(a)**: Error signal that the **MTS** lock generates for an atomic transition of frequency ω_a . Note that the background signal is flat, which eliminates the possibility of unintentionally locking to a different, nearby resonance. **(b)**: Electronics for the cooling laser's **MTS** lock, as described in the main text.

phase difference, and therefore the larger the error signal, which follows equation 5.26.

The **FWM** process naturally eliminates the effect of Doppler shifts on our error signal, because **FWM** and therefore the modulation transfer can only occur when both the pump and probe are simultaneously resonant with the atomic transition: and here, they are travelling in opposite directions. Furthermore, in comparison to frequency locking with pump-probe spectroscopy [164]⁴⁴, the **MTS** error signal has a flat background, because the lineshape and therefore $F(\omega)$ only includes the absorption peak for the cooling transition and not the Doppler-broadened lineshape on each side of it. Also, it is only significant for the cooling transition, because the cooling transition is the only cycling transition.

Once we have the error signal, we can then lock the laser. Figure 5.9b shows the electronics we use for this. By inputting both the **EOM** drive signal and the photodiode probe beam signal to a frequency mixer, we can get a beat signal. The mixer outputs the sum and difference frequencies of the two inputs: with a phase shifter and a low-pass filter, we can isolate just the difference/beat signal, which is indeed our error signal. We then regulate the error signal using a proportional-integral-derivative (PID) controller, and finally feed back the regulated

⁴⁴ Sometimes referred to as saturated absorption spectroscopy.

signal to the laser cavity piezos and laser current in order to change the laser frequency. Thus, we now have a feedback loop, keeping the laser locked on the cooling transition.

Repump laser lock

For the repump laser, we use a beat-note offset lock [165]. The cooling laser is already locked via MTS, so we can use that as a stable reference and offset its frequency by 6.568 GHz (see figure 5.8) to lock to the repump transition. Thus, we transfer the stability of the cooling laser onto the repump laser.

Figure 5.10 shows how we implement the offset lock. The method is as follows: we tap off a small amount of cooling and repump light, and interfere these two beams on two input ports of a beamsplitter. This results in a beat-note at frequency $\omega_{\text{beat}} = |\omega_{\text{cool}} - \omega_{\text{repump}}|$, which we detect on one output port using a fast photodiode.⁴⁵ Then, we mix the beat-note with a signal from a reference oscillator at frequency ω_{LO} . Next, we use a method akin to balanced detection: we split the mixer output in half, and send to two detectors labelled A and B, which have a low-pass and a high-pass filter respectively. If we choose the filter cutoff frequencies such that they have a crossover region, then in that region there will be precisely one frequency ω_{locked} where both A and B detect the same signal magnitude. If the laser is erroneously lower in frequency, A will detect a greater signal; if higher, then B will detect a greater signal. Therefore, if we tune our reference oscillator such that $\omega_{\text{locked}} = |\omega_{\text{beat}} - \omega_{\text{LO}}|$, then the difference in the signals measured by A and B directly corresponds to the error signal. Finally, we regulate the error signal and feed back to the repump laser frequency to lock it.

Note that in figure 5.1, we have an single-pass AOM in the repump beam line so that we may quickly turn it on/off. This AOM will detune the beam. Therefore, we adjust ω_{LO} slightly so that the beam after the AOM is resonant with the repump transition.

⁴⁵ Vescent D2-160 High-Speed Beat Note and Pulse Detector

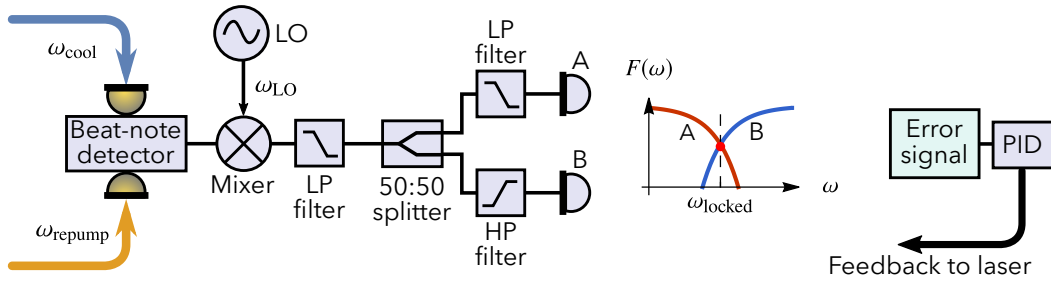


Figure 5.10: Electronics for beat-note offset lock, as described in the main text. Inset shows the filter responses for detectors A and B.

To summarise this section on the experimental design, we designed an atom loading scheme using a MOT placed directly above the cavity as a cold atom source. We have designed it such that we may either use the MOT as-is for probabilistic loading, as we will describe in section 5.6, or combine it with optical tweezers for deterministic loading, which we have allowed for in our design and which is ready to implement.

5.3 Cavity stabilisation

Our cavity has unique stability properties, requiring a unique stabilisation method. In this section, we shall describe these, and how they affect the experiments we perform.

5.3.1 Mechanical stability of the cavity

We desire to lock the cavity, just as we did with the lasers in section 5.2.4, so that we can constantly drive the atom-cavity system at our desired frequency. In particular, we want to lock the cavity to or near to the cooling transition, since using a cycling transition will generate much more fluorescence, because the transition will cycle many times over during the cavity transit time.

How stable must the cavity be to allow us to lock it? We can calculate the required length stability to an order of magnitude as follows. Let us set the condition that the cavity frequency

must not drift by $> \Delta\omega_{\text{FWHM}}/2$ in order to keep it locked. We know the locking frequency must satisfy the standing-wave resonance condition (equation 2.6), so we can therefore relate the frequency drift to equation 2.6 to get

$$\omega_{\text{unlock}} = \frac{\Delta\omega_{\text{FWHM}}}{2} = \frac{n\pi c_0}{L_{\text{cav}}} - \frac{n\pi c_0}{L_{\text{cav}} + \Delta L}, \quad (5.27)$$

where ΔL is the cavity length change past which the cavity will unlock, which we then rearrange for:

$$\Delta L = \frac{L_{\text{cav}}^2 \Delta\omega_{\text{FWHM}}}{2\pi n c_0 - L_{\text{cav}} \Delta\omega_{\text{FWHM}}}. \quad (5.28)$$

Inputting the values from table 4.3, we therefore estimate that we must stabilise the cavity length to ~ 10 pm.

To achieve this length stability, we can use the Pound-Drever-Hall (PDH) technique [161]—the standard method of locking high-finesse cavities. For this method, we use exactly the same frequency modulation as we had previously discussed in section 5.2.4, but with the rubidium atoms replaced by the optical cavity. Instead of looking at the transmitted signal through the atoms, we look at the reflected signal from the cavity. This subtlety makes a significant improvement: we use the transfer function $F(\omega)$ for the cavity, but now, we effectively measure the phase difference between the field directly reflected from the cavity and the field that enters the cavity, circulates, and later leaks back out. This leakage field contains the ‘memory’ of all the phase perturbations that occurred while it circulated, equivalent to \mathcal{F} bounces between the mirrors before escaping. Therefore, this rolling average allows us to measure the sidebands’ phase even very far off resonance, and so we can greatly improve the lock’s capture range by using a larger ω_m , with no extra cost in feedback speed.

We attempted to PDH-lock the cavity by generating this PDH error signal and feeding it back to the cavity piezo(s) to stabilise the cavity length. However, as we shall now see, the

mechanical stability of this cavity is still not sufficient even with the increased capture range afforded by [PDH](#) locking.

In the following two figures [5.11](#) and [5.12](#), we demonstrate the mechanical properties of the cavity.

In figure [5.11](#), we show the power spectrum of the cavity's mechanical vibrations. Here, we turn off the cavity scan and only apply a constant (offset) voltage to one piezo⁴⁶ in order to keep the cavity roughly on resonance: ambient vibrations then vibrate the cavity around and across resonance. We then take a long measurement capturing many resonances, and take the Fourier transform of that to get the power spectrum. The vibration response of the cavity is stronger at 1.2 kHz and 3.3 kHz; these are mechanical resonances. Because these resonances fall within the range of typical acoustic background noise in the laboratory, it becomes much more difficult to frequency-lock the cavity (with the [PDH](#) method), and this is why those attempts were unsuccessful.

In figure [5.12](#), we further isolate the effect of the vibrations on the cavity by measuring its mechanical transfer profile: a function $F(\nu)$ that tells us what vibrations are picked up at frequency ν . Instead of varying the laser frequency as we would do to measure the cavity's optical transfer profile in section [5.2.4](#), we sweep (gradually increase) the frequency of one scanning piezo.⁴⁷ Consequently, the scanning piezo acts as a driver of mechanical vibrations. Moreover, we use a 1064 nm laser: the cavity mirrors are only weakly reflective at this wavelength, and so the detected cavity transmission is the sum of a large constant ([DC](#)) component and a small sine component. We supply the sweep signal using an arbitrary waveform generator ([AWG](#))⁴⁸ so that we get a precise, controllable sweep and so we have a phase reference: by comparing the [AWG](#) phase to the photodiode signal phase, we can

⁴⁶ The other piezo is not used here.

⁴⁷ We use a sine wave for this scan. Normally for cavity scans, we use a triangle wave (see e. g. figure [4.23](#)), but a triangle wave contains components of many frequencies and so is not suitable here.

⁴⁸ Keysight 33612A




Figure 5.11: Power spectrum of the cavity transmission in vacuum without scanning, under ambient conditions. First, we manually apply a voltage offset to one of the two piezos to keep the cavity on resonance with the laser. Then, ambient sound in the laboratory vibrates the cavity around resonance. While this occurs, we take a long measurement of the cavity transmission (oscilloscope trace), recording many resonances. The Fourier transform of that free-running trace gives this power spectrum. We take a few traces at different times, so that we may detect any change in ambient conditions.

therefore see the phase response of the cavity to the mechanical vibrations. To perform this measurement, we send the transmission signal to a spectrum analyser measuring $F(\nu)$ as a signal gain, i. e. ratio of (input) AWG drive amplitude and (output) photodiode signal of transmission, and which also measures the phase between the two.

Note that in both figures 5.11 and 5.12, we took the data after damping the vacuum chamber apparatus with sorbothane sheets between it and the optical table on which it rests.

These pronounced resonances make PDH locking impossible because the capture range and bandwidth of the lock is not sufficient to counteract the large shifts in cavity length. As soon as the cavity encounters a mechanical resonance, the length will shift. Is it possible to counteract this? Consider the shape of the phase signal in figure 5.12. From section 5.2.4, we know that ϵ is proportional to the phase difference from the reference (i. e. from resonance),

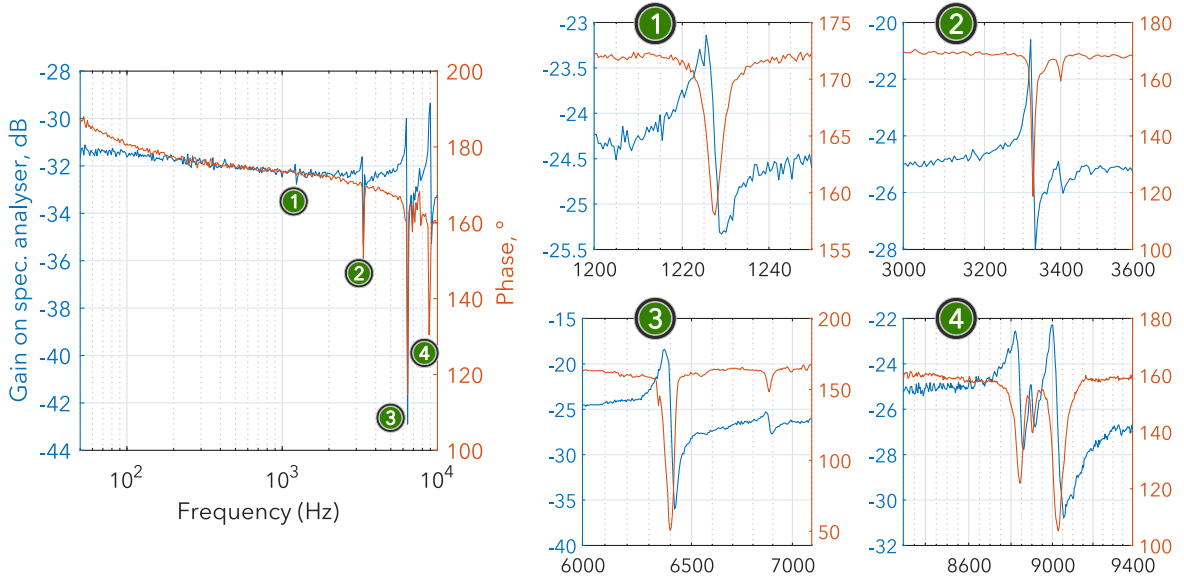


Figure 5.12: Mechanical vibration response/transfer profile of the FFPC in vacuum. Each numbered plot zooms in on a region of the main graph on the left. Here, we input a 1064 nm laser to the cavity. Simultaneously, an AWG drives one piezo with a sinusoidal frequency sweep. The low \mathcal{F} /high transmission with this laser means that the detected transmission (blue) is the sum of a large constant (DC) component and a small sinusoidal component. The second piezo provides an offset to keep the cavity on resonance. Simultaneously, we measure the relative phase of the AWG drive and cavity transmission sine (red).

but our measured phase signal does not have the dispersive profile matching figure 5.9a: the gradient has the same sign on either side of the resonances, so the lock electronics would not know whether to apply a positive or negative feedback voltage to compensate for the error. Instead of counteracting the effect, can we filter out the resonance from our lock electronics? We cannot, because we have mechanical instability in the kHz regime. For slow vibrations ~ 100 Hz, we can apply damping materials and a low-bandwidth lock. For fast vibrations ~ 100 kHz, we can apply a low-pass RF filter. But, in the kHz regime, we cannot apply a sharp enough notch filter, nor can we apply a lock of high enough bandwidth to compensate these large vibrations.

To address this instability, we must first understand its cause. By attaching an accelerometer to various parts of the vacuum chamber to locate where the instability comes from, and by

considering the geometry of the setup, we therefore hypothesise that the cavity mounting finger (see figure 5.7) is the likely source of the instability. Its long, flat structure may be acting like a cantilever, with the cavity mounted at the end of it. If the resonant frequencies/harmonics of the cantilever are close to what we see in the above figures, then we can verify our hypothesis. Firstly, we calculate the harmonics of a long, flat cantilever, not including the coned section, using [166]

$$\omega_n = \left(\left(n - \frac{1}{2} \right) \pi \right)^2 \frac{1}{l} \sqrt{\frac{Y \left(\frac{bh}{12} \right) (b^2 + h^2)}{\rho bh}} \quad (5.29)$$

where ω_n is the oscillation frequency of the n^{th} harmonic, Y is the Young's modulus, ρ is the density and $(l \times b \times h)$ are the dimensions of the cantilever. Using: $Y = 180$ GPa, $\rho = 7400$ kg m⁻³ and $(l \times b \times h) = (97.3 \times 13.0 \times 2.9)$ mm; gives the fundamental harmonic $\omega_1 = 4.9$ kHz, ignoring the coned section. Although this is at the right order of magnitude, the difference from the lowest harmonic of 1.2 kHz resonance suggests that the coned section makes a significant difference to the vibration frequency.

To improve our model, we therefore use finite element analysis (FEM) software⁴⁹ to identify the resonant frequencies of the cavity finger. Figure 5.13 shows the results of the FEM analysis; with a resonance visible at 1.2 kHz and other resonances nearby to those in figure 5.12, we can reasonably expect that the cavity finger is responsible for the instability. We note that more recent designs of mounting fingers for new experiments in our research group have now incorporated such analysis to design stable mounts: firstly by using FEM, and secondly by adding as much mass as possible to increase the resonant frequencies to a point outside the acoustic range where we can more easily damp them.

Given that PDH locking is not possible, how should we stabilise the cavity? For our particular shape of the phase signal, we can use a side-of-fringe lock to keep the resonance at a well-defined reference position in the scan: we can define a phase 'set point' at the centre

⁴⁹ Nastran for Autodesk Inventor




Figure 5.13: Mechanical vibration response of the cavity mounting finger, calculated using [FEM](#). To run this simulation, we apply an oscillating force to the cavity finger at the point closest to the cavity at a range of frequencies from 800 Hz to 9500 kHz, and see by what magnitude the cavity accelerates by (i. e. its vibration response).

of the scan, and if the phase deviates from the set point, we can nudge the cavity position forwards or backwards depending on whether the cavity resonance drifted in front or behind the scan centre (left or right on an oscilloscope trace). In the next section, we therefore describe how we stabilise the cavity in the centre of the scan using a side-of-fringe lock with a lock-in amplifier, allowing us to keep the cavity roughly close to atomic resonance. Thus, we can still take our fluorescence data even in spite of the large mechanical vibrations.

5.3.2 Frequency stabilisation of the cavity

While scanning the [FFPC](#), we use a lock-in amplifier⁵⁰ to compensate for medium- and long-term drifts in the cavity resonant frequency, thus keeping the the resonance in the centre of the scan and allowing us to ‘zoom in’ on it, increasing the interaction time of the cavity with atoms. The lock-in amplifier works similarly to the [MTS](#) frequency mixer electronics in figure 5.9b, where we compare the phase of a reference signal and the signal of interest by beating them and isolating just the [DC](#) signal corresponding to the phase difference. From

⁵⁰ EG&G Princeton Applied Research 5210.

figure 5.9b, we replace the EOM with the cavity-scanning piezo, the rubidium cell with the cavity, and there is no pump beam. Here, we set up the lock-in so that the amplitude of the DC signal, which will serve as our error signal to feed back to the cavity length, corresponds to the phase difference between the centre of the cavity resonance and the centre of the scan.

We set up the stabilisation as follows: a sine signal from the AWG drives one of the cavity piezos and is also sent to the lock-in amplifier. Thus, the AWG is the local oscillator for the lock-in, and will serve as our reference frequency. Meanwhile, we set an offset voltage to the other piezo to keep the cavity roughly on resonance. By sending the cavity transmission photodiode signal to the lock-in, we can then measure the phase difference between the reference and the transmission signal. Now, we tune the phase difference by phase-shifting one of the signals with the lock-in: we set the phase difference to be equal to zero at the transmission peak. Therefore, the lock-in output is the error signal. We then regulate the error signal with a PID controller and feed it back to the scanning piezo to stabilise the cavity length.

With our lock-in stabilisation method, we keep the position of the cavity resonance peak at the same position for many scan cycles, and also in the centre of the scan. Since we know that the cavity mechanically resonates at 1227 Hz, we scan the cavity close to this frequency to increase the amplitude of the resulting error signal and stabilise the cavity better. Figure 5.14 demonstrates how we harness this resonance to improve the lock-in performance.

When operating the lock-in, we found that the minimum scan range of the cavity was still ~ 1 GHz: many times larger than the cavity linewidth $\Delta f_{\text{FWHM}} \approx 20$ MHz. Below this minimum scan range, the cavity became unstable and started to display multiple resonances around the scan centre. At this range, the lock-in cannot correct length perturbations quickly enough, due to the finite time constant of its filter. The time constant effectively sets the




Figure 5.14: Histograms characterising how, when we activate the feedback from a lock-in amplifier, we harness the cavity’s mechanical resonance to mitigate vibrations. If we drive the cavity with a 1 kHz scan, versus with a 1223 Hz closer to the resonance, then we keep the cavity transmission peak better-centred on the scan (left plot) and reduce the jitter between subsequent peaks (right plot). Here, we take jitter to mean deviation from the expected time difference between one transmission peak and the peak on the next scan cycle. For example, since the two groups of occurrences are closer together for 1223 Hz, the jitter is smaller.

sharpness of the filters used in the lock-in electronics.⁵¹ A smaller time constant means that the lock-in can respond more quickly to length perturbations, but also that the error signal contains a broader range of frequencies after filtering (and we do not want to keep the mechanical vibration frequencies). We use a 30 ms time constant, effectively averaging the position over many peaks/cycles of vibration. Hence, we roughly keep the cavity in the same position in successive scan periods. This way, we can expect the cavity to be at or close to the centre of the scan, even under the effect of ambient vibrations in the laboratory. Note that the minimum scan range with the lock-in before encountering instability is still an improvement over using no lock-in (about several GHz).

⁵¹ Literally, the time constant is the reciprocal of the bandwidth of the low-pass filter.

5.4 Experiment methods introduction

We now want to examine the characteristics of our cavity's atom-photon coupling capabilities by fulfilling two key requirements:

1. Make the [FFPC](#) resonant with the desired atomic transition;
2. Load atoms into the cavity (mode).

We are looking for, in the first instance, the Purcell enhancement of atomic emission into the cavity. We will perform two experiments to do this, described in the next sections.

For the case of one or a few atoms, the atom-cavity coupling splits the cavity resonance into two resonances as we discussed in [chapter 4](#), manifesting itself as a sudden decrease in the number of photon counts on a detector that observes the cavity output. This is because, with the atom inside, the coupled atom-cavity system is no longer resonant with the incoupled fibre light. In principle, we can exploit this to measure the atom-cavity coupling strength. If we ensure that we load n atoms into the cavity, then instead we can detune the laser by $g\sqrt{n}$ from the uncoupled cavity resonance and observe a transmission peak as the atom falls in and couples to the cavity. However, because of the mechanical stability issues described in [section 5.3](#), we stabilise the cavity around the cooling transition and scan over it, rather than lock the cavity directly to it.

Once we implement the tweezers into the cavity system, we can use this splitting to definitively trap atoms in the cavity. At the instant we detect the transmission decrease, we would turn on the tweezers and trap the falling atom(s), in much the same way as when intracavity dipole force traps were used to trap atoms falling into macroscopic cavities [[39](#)].

For both of our experiments, one may expect that we directly observe the vacuum Rabi splitting, just like the experiment of [Bernadot et al. \[167\]](#). However, our experimental parameters are different (e. g. atom number in cavity, coupling strength, optical rather than

microwave frequencies, experiment timescale), and so we instead focus on observing the Purcell-enhanced fluorescence of a few loaded atoms only, where we expect to see an increase in the detected signal when the laser-driven atoms couple to the cavity.

5.5 Beam of atoms: pushing atoms into the cavity

For the first experimental configuration, we keep the MOT on, and scan the cavity across the resonance(s) of the coupled system. Our aim is to push atoms from the MOT directly into the cavity using a push beam, giving a steady flux of atoms that transit the cavity mode [168, 169]. The push beam passes directly through the MOT and through the cavity, pushing atoms out of the MOT towards the cavity by radiation pressure. Our scheme and parameters are similar to [168], but otherwise we use a smaller MOT and lower power beam. Accounting for that, we can make an order of magnitude estimate of the atomic flux across the cavity as $\sim 10^4$ atoms s^{-1} . The push beam also drives the atomic transition (section 5.6.2 will describe this function in more detail), and so the atoms will scatter light into the cavity mode, which then escapes through the output fibre. Because the detected signal is small, we will use a single-photon counting module (SPCM)⁵² to monitor the cavity output.

5.5.1 Results of atom-pushing experiment

Figure 5.15 shows histograms of the SPCM counts when the MOT is on and off. The cooling and repump resonances are on the left and right respectively in all three histograms. To take ‘MOT off’, firstly we switch off the MOT and use the lock beam—the beam used for the lock-in stabilisation—to quickly find the resonances while scanning the cavity over a ~ 10 GHz range featuring the cooling and repump transitions. Then, we block the lock beam

⁵² Excelitas SPCM-780-14-FC




Figure 5.15: Histograms that indicate polarisation drift of the cavity-collected fluorescence from a push beam. The cooling and repump resonances are on the left and right respectively in all three histograms. First we take a measurement with the **MOT** off, then immediately after we take a measurement with the **MOT** on. We also repeat this, and take a second measurement with the **MOT** on. The fibre polarisation drifts differently for these two histograms, resulting in different proportions of counts for the two resonances and therefore indicating atom-cavity interactions.

and look for the resonances without it. Atoms from the background gas, being driven by the push beam, fluoresce into the cavity mode. We pick up that signal on the **SPCM**, and the resulting histogram of that signal over the order of minutes gives ‘**MOT** off’. Immediately after, we take ‘**MOT** on’ by switching on the **MOT** and repeating the measurement. We repeat the whole process twice, to get two ‘**MOT** on’ histograms. The cavity drifts position over the minutes it takes to make the new histogram. However, at the same time, the polarisation of the **FFPC**’s fibres drifts. Thus, for the two ‘**MOT** on’ measurements, we have different counts for the cooling and repump resonances. In this way, even though fluorescence of the background gas accounts for the majority of counts, we can attribute the polarisation-dependent shift to Purcell-enhanced emission of the flux of atoms from the **MOT**, since those are in a particular polarisation state whilst the background gas atoms are unpolarised.

We desire to investigate the polarisation dependence further. Since the lock light’s polarisation can be rotated before its input to the **FFPC**, we could in principle separate the lock light from the scattered light by using different polarisations for them. However, the extinction ratio of our polarisation filters ($\sim 10^{-3}$) and the fact that the lock light is $\gtrsim 10^3$ more powerful than the scattered light makes this impractical. We also cannot detune these

two beams far enough from each other with our [AOMs](#) so that we filter them out by frequency. For this reason, in the next experiment we will separate locking light and driving light by switching on and off these beams with the [AOMs](#).

Note that for the measurements where the [MOT](#) stays switched on, the magnetic field from the [MOT](#) coils will Zeeman-shift the cooling and repump frequencies slightly. Using $\Delta f = \Delta E/h = \mu_B B_z/h$ [147] where h and μ_B are Planck's constant and the Bohr magneton respectively, and the field gradient from section 5.2.2 and height above the cavity of 1 mm, an order of magnitude calculation yields $\Delta f \sim 1$ MHz. (We approximated the Lande- g factors as unity for both the cooling and repump transitions). So, we neglect the effect.

5.6 Probabilistic loading: dropping atoms through the cavity

For the second experimental configuration, we use a probabilistic loading scheme to load a few atoms into the cavity at a well-defined time, and scan the cavity across the resonance(s) of the coupled system. Our loading scheme can be summarised in two steps: quickly form a [MOT](#), then turn it off to drop the atoms into the cavity. As figure 5.1 shows, a driving beam from the cavity side causes the atoms to scatter photons into the cavity mode, which then couple into the output fibre and are detected by the [SPCM](#). During the second step, we do not send any light to the cavity, so that lock light does not drown out the fluorescence signal or overload the [SPCM](#).

5.6.1 Loading the cavity with atoms

To load the atoms into the cavity by gravity (rather than relying on the push beam), we position the [MOT](#) directly above it, and then turn it off: thus, the atoms fall through the cavity.

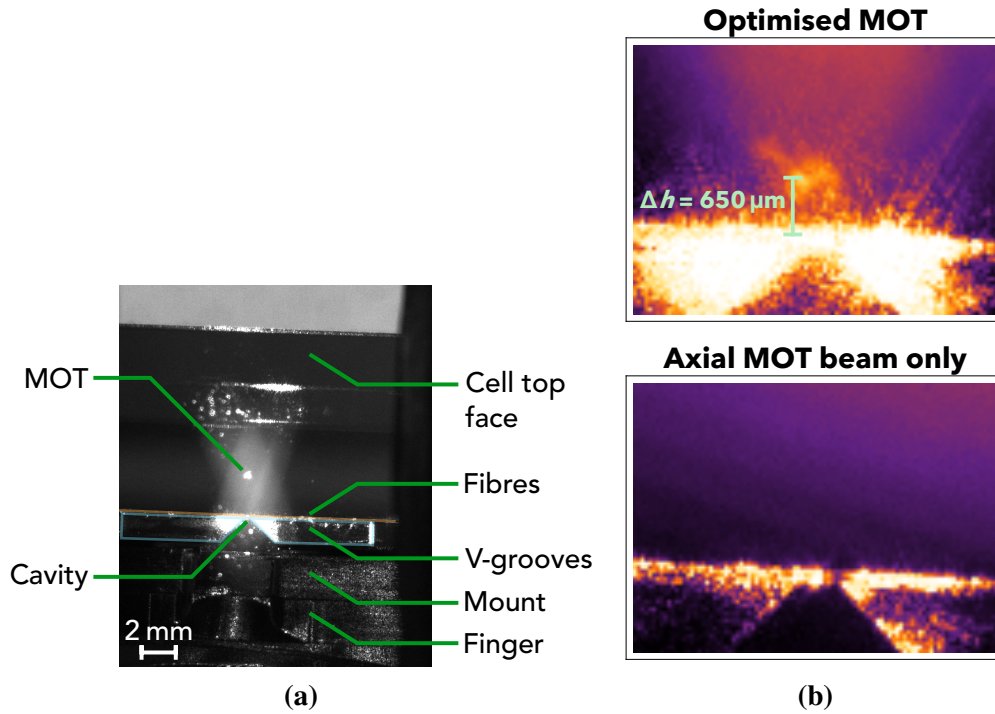


Figure 5.16: IR images of the MOT for the cavity loading scheme. (a) Initial MOT that we formed, clearly showing where it is relative to the other parts of the experimental setup. Here, the MOT is $\Delta h = 2 \text{ mm}$ above the FFPC. (b) Optimised MOT (in false colour). We brought it closer to the cavity for better cavity loading. To more accurately measure its size, we first focus the camera on the cavity and then point it at a calibration rule. By taking an image of the rule in focus, we obtain the size of one pixel. We then measure Δh and the MOT $1/e^2$ width in terms of pixels. We confirm that this is indeed a MOT by blocking the two diagonal MOT beam pairs: the lower image shows that the MOT disappears as a result.

Figure 5.16 shows the two MOTs we used for our scheme. We started with an initial MOT, and then tweaked the beam and coil alignments to obtain a loading-optimised MOT. In our experiments and in the following calculations, we use the loading-optimised MOT.

We want to ensure that, on average, we load only a few atoms into the cavity at a time. If there are too many atoms, e. g. ~ 1000 on average, the collective coupling strength $g\sqrt{n}$ of the atoms would split the cavity resonance by several GHz, too far for our driving laser to reach with our available modulating devices (AOMs). Moreover, collective interactions of

the atoms with different coupling strengths would be difficult to analyse.⁵³ If there are too few atoms, e. g. $\lesssim 1$ on average, then the scattering from background atoms may drown out the Purcell-enhanced scattering from loaded atoms. Therefore, we should have a good idea of how many atoms are in the cavity, so that we can focus on detecting Purcell enhancement.

How many atoms are likely to be loaded in the cavity after we drop the MOT? We can calculate this by assuming the MOT cloud follows a ballistic trajectory upon release. Then, we just account for the expansion of the MOT cloud as it heats up after being released. From this, we can estimate the density of atoms in the cavity. Since we already know the cavity mode volume, therefore we will know the number of atoms loaded in it.

To perform our calculation, firstly we need to know the initial MOT density. Since the MOT beams are Gaussian, we can safely assume that the MOT density itself is also Gaussian. The MOT atom density in 1D is therefore

$$\eta(x) = \frac{N_{\text{MOT}}}{\sigma_x \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{x}{\sigma_x}\right)^2}, \quad (5.30)$$

where x is the position from the centre of the MOT, σ_x is the $1/e^2$ MOT width along x , and which we have normalised such that it contains N_{MOT} atoms in 1D, $\int_{-\infty}^{\infty} \eta(x) dx = N_{\text{MOT}}$. Since our MOT beam widths and powers are all equal, and assuming our MOT is spherical, then the MOT density is the same along each dimension. Therefore, the MOT density in 3D is

$$\eta(r) = N_{\text{MOT}} \eta_x(x) \eta_y(y) \eta_z(z) = \frac{N_{\text{MOT}}}{(2\pi)^{\frac{3}{2}} \sigma^3} e^{-\frac{1}{2} \left(\frac{r}{\sigma}\right)^2}, \quad (5.31)$$

where r is the radial distance from the MOT centre and we have now normalised η such that it contains N_{MOT} atoms in 3D.

⁵³ Although we would not expect to see superradiance because $\lambda \ll L_{\text{cav}}$ [170], we may see a large spread of detection times due to collective absorption/dispersion, somewhat similarly to electromagnetically induced transparency (EIT) [171].

Next, we need to calculate how the MOT expands once released. Let us release it at time $t = 0$ by switching off the beams and turning off the coils. Then, using equation 5.12 and assuming it expands equally in all directions from a central point, we get at time t

$$\sigma^2(t) = \sigma^2(0) + \frac{k_B T}{m} t^2, \quad (5.32)$$

where T is the MOT temperature and m is the mass of a single ^{87}Rb atom. We now want to know what the MOT density is when it reaches the cavity. Using the standard equations of motion under constant acceleration, we get the simple expression $t^2 = 2\Delta h/g$, where Δh is the height of the MOT above the FFPC and g is the acceleration due to gravity. We input this to equation 5.32 to get

$$\sigma_{\text{arrive}} = \sqrt{\sigma^2(0) + \frac{k_B T}{m} \frac{2\Delta h}{g}}. \quad (5.33)$$

Substituting this into equation 5.31 allows us to calculate the peak atom density η_{peak} at the centre of the cavity as the MOT arrives,

$$\eta_{\text{peak}} = \frac{N_{\text{MOT}}}{(2\pi)^{\frac{3}{2}} \sigma_{\text{arrive}}^3}, \quad (5.34)$$

and therefore estimate the maximum number of atoms loaded into the cavity after one MOT fall,

$$N_{\text{cav,max}} \approx \eta_{\text{peak}} V_m, \quad (5.35)$$

where V_m is the cavity mode volume.⁵⁴

We can now estimate the number of atoms in the cavity. Using equations 5.31 and 5.35, figure 5.16 and the following values: $\eta(r = 0, t = 0) \sim 1 \times 10^{17} \text{ m}^{-3}$ [150, 153], $\sigma_0 = 300 \mu\text{m}$, $m = 1.4 \times 10^{-25} \text{ kg}$, $\Delta h = 650 \mu\text{m}$, $T \sim 100 \mu\text{K}$ ⁵⁵ and $V_m = 4 \times 10^{-15} \text{ m}^3$; we have

⁵⁴ We assume $V_m \ll \sigma_{\text{arrive}}^3$.

⁵⁵ This order of magnitude estimate agrees with the measured temperatures of similar MOTs in our research

$N_{\text{cav,max}} \sim 7$ atoms on average. Therefore, we can assume that for our experiments, there are a few atoms in the cavity, which is what we want to detect Purcell enhancement.

As for when we expect to observe the fluorescence after releasing the MOT, the expected travel time to the FFPC is $t = 11$ ms. The expected interaction time, which we assume to be equal to the transit time across the mode, is therefore $\sim 90 \mu\text{s}$:⁵⁶ from the value of $1/\Gamma$, we can therefore drive a transition many times over before the atom exits the cavity.

5.6.2 Driving beam

For this experiment, we use the push beam to drive the cooling transition for atoms in the cavity. As mentioned before, this beam has dual functionality as both a push and driving beam. But now, we focus only on the driving effect, and henceforth we rename it as the driving beam. We expect to measure an increased fluorescence signal even from a few atoms, because as the atoms excited by the laser relax to the ground state, a large fraction of their emission will be into the cavity mode due to Purcell enhancement. We set the driving beam intensity at a few times I_{sat} for the D_2 transition to ensure the emission signal we detect is larger (as the atom will spend longer on average in the excited state). However, note that the cavity is resonant for only $1/500^{\text{th}}$ of the scan period (when using the lock-in).

We must highlight a subtlety: the driving beam and lock beam are close by in frequency. The driving light is necessarily resonant with the cooling transition, while the lock beam is Doppler-detuned by 2.25Γ (as it is derived from the MOT cooling beam). Since the cavity scan, at ~ 1 GHz, is large compared to 2.25Γ , this does not affect our results. Also, given that the time constant of the lock is 30 ms, and the length of a launch sequence is 25 ms, the lock-in stabilisation is therefore still effectively active during the data acquisition.

group [153].

⁵⁶ These are both calculated using the standard equations of motion for constant acceleration.

5.6.3 Atom loading sequence

We want to repeat our cavity loading many times to take more data, and take this data only when necessary: that is, when we drop the MOT through the cavity. So, we use a computer-controlled loading sequence to turn on and off the relevant parts of the experiment. Figure 5.17 shows the sequence timings. From left to right: firstly, we start scanning the cavity length at a frequency of 1 kHz and amplitude corresponding to ~ 1 GHz. Then, we trigger our oscilloscope 0.5 ms later.⁵⁷ Afterwards, we switch the MOT off in two stages: firstly we switch off the current to the MOT coils, and secondly we switch off the cooling beams by stopping driving the AOM⁵⁸ used to detune them. We do this because switching off the coils takes about 2 ms [155],⁵⁹ whereas switching off the cooling beams via the AOM takes only 20 ns. Thus, the molasses keep the atoms cool until the moment we turn them off and therefore release the MOT. The locking light is part of the same beam path as the cooling light, and so switches off at the same time. Once the atoms reach the cavity, they couple to the cavity mode, and the driving beam causes them to scatter photons into it. We then detect these photons using the SPCM on the FFPC output fibre. We take data only during this period while the lock light is off, so that we only look at detector clicks resulting from the driving light. In total, the loading sequence lasts 25 ms.

At the end of the loading sequence, we turn the MOT back on and wait for 1.0 s before starting the sequence again, in order to repopulate it. We repeat the sequence many times to acquire the data shown later in section 5.6.4. Throughout, the drive beam remains on permanently.⁶⁰

⁵⁷ The delay is equal to half a cavity scan cycle, so that detection starts in the centre of the scan—where the resonance is, and where we want to look for detector counts.

⁵⁸ Gooch & Housego AOMO 3110-120

⁵⁹ We use one third of the current that Dilley used, who measured a 5 ms switching time.

⁶⁰ The repump beams for the MOT are also on permanently.




Figure 5.17: Timings of the cavity loading sequence. We turn off the locking light and MOT cooling beams with an AOM. We trigger our oscilloscope slightly after we trigger the scan, so that they are synchronised properly. The repump and driving beams remain on permanently. Every 1 s, we repeat this sequence.

5.6.4 Results of probabilistic loading

Figure 5.18 shows a set of histograms that demonstrate how the presence of the cavity Purcell-enhances the emission of the laser-driven atoms that transit the cavity mode. The cavity effectively channels the emission into the output fibre, where we detect the fluorescence on the SPCM. As we mentioned prior, the cavity was scanning while we observed this fluorescence: therefore, we have many peaks in the detector counts across the 20 ms data-taking period. To gather enough counts for the histograms, we repeated the launch sequence described in section 5.6.3 500 times (except for ‘molasses off’, which we refer to soon). We wrap all of these peaks into a histogram for a single 1 ms scan cycle by overlapping all of the cycles. We have two peaks, since the cavity is resonant twice during a cycle: once on the piezo upstroke, and once on the downstroke.⁶¹ Since we wrap the peaks, we take the mean counts over all the cycles and use that to calculate the standard error, rather than the usual Poissonian \sqrt{N}

⁶¹ Just as in figure 4.23, we do this to ensure that the observed resonances are not spurious or changing significantly during scans.

uncertainty for counting single photons [172]. We apply background subtraction to these plots; we see in the plot of the error intervals that the negative probability densities are within a positive value for their error, so we can be confident that the errors describe the data correctly and are physical.

Now, we are interested in the data where the MOT was on and we executed the experimental sequence. However, the fluorescence we observed may be due to spurious sources. For example, background atoms that did not arrive from the MOT may have transited the cavity, generating fluorescence. Therefore, we repeated the fluorescence experiment under four different experimental conditions corresponding to the four histograms:

- ‘MOT on’: Execute the normal experimental sequence where we drop the MOT through the cavity.
- ‘MOT off’: Do not turn on the MOT coils, otherwise executing the normal sequence.
- ‘Molasses off’: Block the MOT beams at their input fibres, and do not turn on the MOT coils. We do not execute the sequence: we continuously take data.
- ‘Dispenser off’: Switch off the rubidium dispenser and let the ion pump completely evacuate any background rubidium before executing the sequence.⁶²

Let us discuss each case. For ‘dispenser off’, there are effectively no rubidium atoms in the chamber, so we do not observe a meaningful signal. Thus, we verify that the beams are properly aligned and do not scatter light directly into the cavity. For the remaining three histograms, we do observe resonances. Since during the measurement the only light reaching the cavity is the driving beam striking through the open side, then these resonances can only be due to atom-cavity coupling involving atoms from the background gas and atoms falling from the MOT that traverse the cavity mode. Moreover, when the cavity is not

⁶² We wait about one day for the pressure to fall adequately.

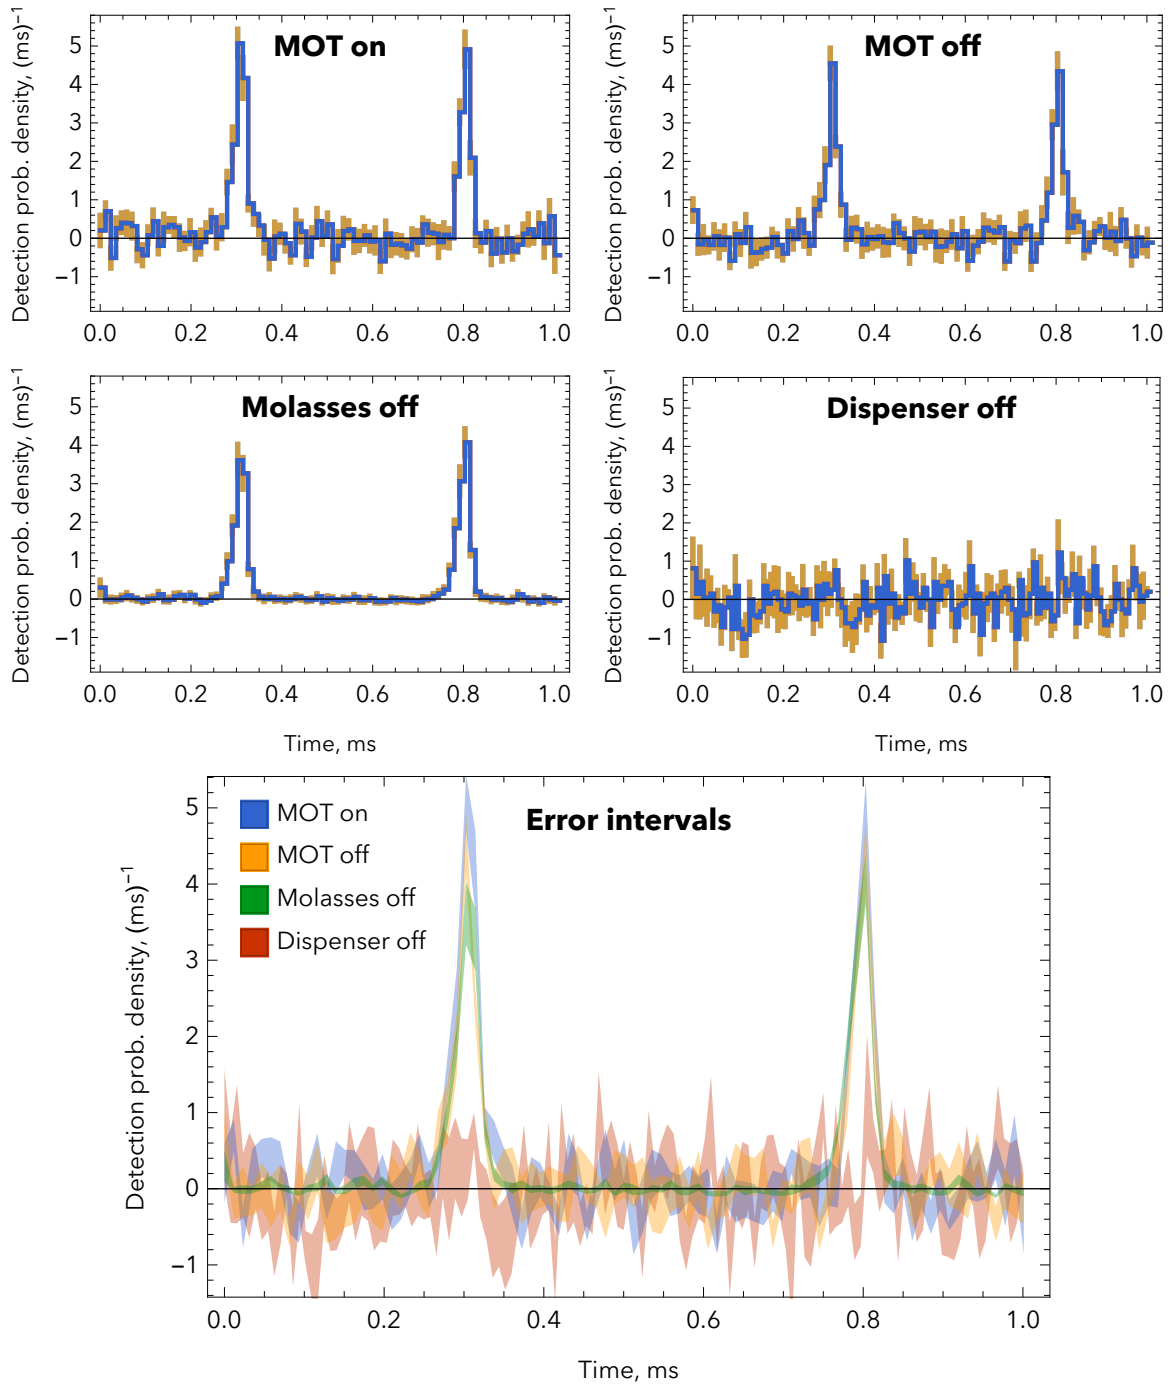


Figure 5.18: Histograms of fluorescence counts showing cavity Purcell enhancement. Each histogram shows the probability density of a detection (SPCM click) occurring at a certain time in a single scan cycle of the FFPC (blue), and error bars for each histogram bin (orange). Each of the four configurations are further described in the main text. For the same reasons as in figure 4.23, we show the peak on both the upstroke and the same peak on the downstroke of the scan. Error bars are standard errors (standard deviation of mean) over many cycles. The lower plot shows only the errors of all four plots together and represented as interval bands rather than bars, for direct comparison of the errors.

resonant/coupled to the atoms, the fluorescence is suppressed to the background level of ‘dispenser off’. Therefore, the cavity must be Purcell-enhancing the emission of these atoms.

Let us compare the relative differences in counts. We note from ‘molasses off’ that fluorescence from the background rubidium is responsible for a large proportion of the total counts: it peaks at roughly 4.0 counts per millisecond on resonance compared to 5.0 counts per millisecond for ‘MOT on’. For ‘molasses off’, we were able to take data continuously since we do not need to execute the launch sequence: therefore, its uncertainties are smaller than for the other plots. By comparing the ‘molasses off’ error bars with ‘MOT on’ and ‘MOT off’, we are confident that the atoms loaded from the MOT are responsible for the increase in detected fluorescence and can rule out other sources.

Let us compare ‘MOT on’ with ‘MOT off’. We see the former has about 0.5 more counts per millisecond on resonance. However, their error bars overlap slightly. This suggests that the loading efficiency from the MOT is low; further adjustments to the MOT alignment should allow us to increase this efficiency and eliminate any overlap of error bars. Furthermore, we note that since the driving beam also acts as a push beam, then the transit time of atoms through the cavity may be reduced as the push beam accelerates the atoms through the cavity, reducing the counts. However, we are already on I_{sat} with the driving beam, so reducing it further will also reduce the fluorescence we measure. So in this case, it is better to leave the driving beam at I_{sat} and note the possible systematic reduction of counts due to this. To investigate this further, we can instead change the MOT density by changing the coil current: the MOT trapping force should scale linearly to this, and we should see the counts scale linearly too [150, 173]. However, for our tight alignment near the cavity, we do not have much leeway in changing the current before the MOT disappears.

We notice a broadening of the resonances. Even though the cavity is only resonant with the cooling transition (and therefore roughly resonant with the atom) for $1/500^{\text{th}}$ of the scan

period, the observed peaks take up a much larger proportion of the scan. Especially for a cavity in the strong coupling regime, Purcell-enhanced emission may be broadened [174]. The magnitude of the broadening depends on the coupling strength. In future, we could use this broadening to quantify the coupling strength by comparing the emission linewidth of the coupled system to a system that we decouple from the cavity by laser-driving the $|F=2\rangle \rightarrow |F'=1\rangle$ transition instead (which, according to figure 5.8, would be detuned by 0.4 GHz from the cavity). However, we would observe a much lower rate of scattered photons from the decoupled system, so we would likely need to stabilise the cavity and implement a PDH lock before attempting to quantify the effective coupling strength.

To summarise, our measurements show Purcell-enhanced emission of the transiting atoms into the cavity mode, even for our scanning (rather than locked) cavity. Given that the cavity is only resonant with the atom for $1/500^{\text{th}}$ of the scan period, this indicates our cavity's coupling strength is likely strong, since it would be otherwise difficult to observe the emission from a few atoms without the Purcell enhancement from the cavity. Since a significant proportion of the emission is due to atoms from the background gas, in future we can increase the loading efficiency of the atoms coming from the MOT by further optimising the MOT parameters and also reducing the rubidium dispensation rate for this particular experiment. Of course, by implementing the deterministic loading with the tweezers as we envisage, we will also alleviate most of the above considerations.

5.7 Conclusion

In this chapter, we have firstly designed our atoms-in-tweezers experiment, with the tweezers themselves ready to be implemented in the cavity, allowing for fully deterministic and reconfigurable trapping of multiple atoms in the strong coupling regime. Combining the state

control afforded by the cavity coupling with the moving and addressing of single atoms will be a significant advancement in the field of quantum information science, and our design sets a clear path to doing so.

Secondly, progressing with our cavity implementation, we have performed our first atom-cavity experiments using probabilistic atom loading. We dropped atoms from a MOT through the cavity under gravity, and measured the Purcell-enhanced fluorescence into the cavity mode while laser-driving the atoms from the open side.

However, we found that the mechanical instability of our apparatus prevents us from performing these experiments in the usual manner. While we deduce that this is not due to the cavity, we still must address the instability before we can perform our envisaged tweezer experiment. Hence, instead of the standard PDH locking technique, we used a lock-in amplifier to provide feedback to stabilise the cavity position while scanning the cavity length. Meanwhile, we suggested how to improve the mechanical stability, and since we have designed the cavity mounting finger to be replaceable, we can adjust its design to eliminate any stability issues in future setups. Some of these adjustments are currently being incorporated in the next iteration of atom-cavity experiments in our research group.

When we performed our experiment, the mechanical instability meant we had to scan the cavity and take a histogram of the fluorescence (detector counts) over many loading sequences. Combined with the probabilistic loading, the fluorescence signal from only the MOT-loaded atoms becomes more difficult to detect. Nevertheless, we did detect it, observing the enhancement by the cavity as a broadened peak of emission events. In future, we can further optimise the cavity loading to investigate further the Purcell broadening effect, and use it to gain a quantitative measure of the atom-cavity coupling strength and indeed verify the suitability of our cavity as an atom-photon quantum interface, as we envisaged in this thesis.

Perspectives on this work

From the work we presented in the previous chapters, we have made progress towards new implementations of fibre cavity experiments and related devices. To give some examples, we:

- optimised atom-photon interfacing with the fibre cavity through simulations;
- devised an assembly method that we subsequently used to assemble a compact fibre cavity with two single-mode fibres with sub-micron precision;
- designed an experiment where we could use the cavity with a reconfigurable 2D array of optical tweezers.

What we do in this chapter is to thread these together, and also thread them with other research. Thereby, we offer perspectives on our work, how it integrates with the literature, and what it allows us to achieve in the future. This chapter roughly resembles an outlook section in a research paper, but with emphasis on the perspectives that our work offers today.

Firstly, we compare our atom-based implementation of cavity networks with other implementations in section 6.1. Secondly, we determine what possibilities there are with our fibre cavity in section 6.2, where we discuss research in the field and then combine our advances with it.

6.1 Other implementations of quantum networks

Neutral atoms and [FFPCs](#) are not the only way to implement the quantum network we envisioned in chapter 1. In addition to the review by Reiserer and Rempe [21], here we discuss some of the matter qubit types which could be used to achieve the same vision (presented as a non-exhaustive list):

- *Ions*: interactions between charged ions due to electrostatic repulsion are much stronger than between neutral atoms. By controlling these interactions, ion trap quantum computers can directly perform high-fidelity logic between the ions [13, 175–177]—something which is more difficult with neutral atoms [178]. This same property also makes trapping ions easier, e. g. using electric fields.¹ Trap lifetimes are at the order of hours [179],² and sideband cooling techniques provide a simple route to increasing the coherence of state control [181]. Even without cavities, creating entanglement for use in quantum networks can be fairly efficient [67]. State interaction is otherwise very similar to atoms. With these benefits, why not favour ions over atoms? The major limitation is charge build-up on the dielectric mirror surfaces of the optical cavity disturbing the electric fields that trap the ions [182]. Therefore, ion-photon interactions in the strong coupling regime become much harder to attain due to the larger mode volume required for the Paul trap [66, 109, 183].
- *Quantum dots*: under perfect conditions, quantum dots could replace atoms as emitters, with the added bonus of their natural protection from decoherence due to their solid-state environment [69]. Quantum dots are specially grown solid-state structures, usually consisting of a nanoscale semiconductor embedded inside a larger ‘host’ semiconductor. A dot’s nanoscale size is smaller than the de Broglie wavelength of

¹ Paul traps are most often used when trapping ions [113].

² Coherence times, however, are still at the order of minutes [180].

its constituent electrons, effectively yielding a single electron hole or vacancy at a well-defined position [184–187]. The potential of each hole is small enough not to mix with neighbouring hole potentials. We can therefore imagine a quantum dot as a zero-dimensional harmonic oscillator. Their level structure is similar to an N -level atom, although it is more difficult to implement e.g. a Raman-based photon production scheme for reasons outlined below. With their much shorter decay lifetime than atoms [147], they are typically used to generate single photons at rapid rates. To use them as quantum memories, we can encode the logical qubit in the spin of the electron involved in the transition. The major disadvantage with quantum dots is inhomogeneous broadening due to their environment: the position of the dots is random, due to the random semiconductor growth process used, so each quantum dot has slightly different properties. If many dots are simultaneously addressed, they need tuning individually. Since Nature defines atoms' properties, they do not need this kind of tuning. Another issue is the spin qubit's interactions with the thousands of nearby nuclear spins in the structure [188], causing rapid decoherence, especially compared to our atomic qubit.³

- *Colour centres*: these consist of a nitrogen- or silicon-vacancy in a nanodiamond, giving an energy level structure that can roughly be addressed in a similar way to a narrow bandwidth atomic energy level.⁴ Due to their solid-state nature, they share the same benefits as quantum dots, with the added advantage that a group of colour centres can be produced more uniformly, with less variation in the properties of an ensemble than quantum dots [64, 65, 190]. Just as with quantum dots, their fundamental disadvantage compared with atoms is inhomogeneous broadening caused by their environment: every colour centre has slightly different properties. They are also not

³ A nuclear spin is roughly 2000 times weaker than an electron spin, but for thousands of nuclear spins their effect will be significant.

⁴ Colour centres are thus closely compatible with telecom wavelengths [189].

being completely isolated from e. g. stray electrons in the diamond. This places a lower bound on the effective linewidth of GHz (when exclusively addressing the colour centre's zero-phonon transition line) [191–193],⁵ and hence the Purcell enhancement by the cavity must be stronger than compared with atoms for an efficient interface [112].

Knowing all of these possibilities for cavity-based quantum networking, we can place our own FFPC device and implementation in better context. Many of the above implementations excel at a particular characteristic. Where our particular fibre cavity excels is with the connection between atoms and photons: because it has good optical access to the cavity mode, and because we have shown in our design and through experiments that it is highly compatible with deterministic trapping using optical tweezers, our device therefore forms a good testbed for investigating the physics of interacting atoms and photons. We stress that our scheme allows us to control the individual atom-cavity interactions, unlike other loosely related cavities where the atoms couple to the evanescent field of tapered nanofibres [195]. Furthermore, our fibre cavity has a predicted atom-cavity co-operativity of 94 ± 5 : well into the strong coupling regime, nearly two orders of magnitude better than other macroscopic cavities created in our research group [55] and double that of the best performing fibre cavities of similar type [174].⁶ In terms of gathering data, the strong coupling and deterministic trapping means we would collect lots of data at a quick rate for the full implementation, similar or better than what Gallego et al. report for their fibre cavity [174]. As a next step forward, with the ability to manipulate individual atoms we could therefore use this as a platform to investigate new quantum technologies with atom arrays [196, 197].

Previously in chapter 5, we had also mentioned future-proofing our experiment design by using a second high-NA lens. This would give us the ability to combine multiple tweezer potential landscapes. How could we use this? For example, using two separate laser beams

⁵ Without this broadening, the natural linewidth of a nitrogen colour centre in diamond is 13 MHz [194].

⁶ See table 4.3 for a full summary of the cavity parameters.

for each lens, we move atoms much longer distances by finely controlling the ratios of the two laser powers to ‘swing’ between the two potentials much more smoothly than can be achieved by changing the **SLM** pattern alone (since the **SLM** pattern can only be modified in discrete steps) [106, 108]. Smoothly modulating the potentials reduces heating of trapped atoms. Similar methods without using an **SLM**/holograms, instead using a type of **AOM** to smoothly move the whole beam, have already been implemented with **FFPCs** [140]. We can further combine these techniques with intracavity trapping precisely at the cavity antinodes [198] and precision alignment of the lenses and cavity [199] to further extend our possibilities. We see that our own cavity, with its good optical access, offers a rich platform for investigating the coupling of atoms that are subject to these complex potentials.

6.2 Present and future fibre cavity research

Since its inception, the research field of **FFPCs** has now become mature enough to catch up with the advances made with regular cavities for use in **CQED**. While regular cavity-based devices have reached the stage where logic between atoms in the same cavity [178] and long distance networking has become possible [24, 200]. Meanwhile, fibre cavity implementations have made steps towards real-world compliance [201], also using heralding and nondestructive detection to make devices resilient to loss [27, 29, 110]. The latter is particularly interesting, as it shows today how fibre cavities are already providing a platform to implement novel and modern quantum technologies. Fibre cavities, being newer and being at the microscale, are potentially more difficult to manufacture, which is why we have not necessarily seen the convergence of standard cavity and fibre cavity experiments yet. But we can expect to see this convergence, for the reasons below: and so we can espouse the benefits of our own fibre cavity.

Consider the history of fibre cavity research. Apart from the advantages of fibre cavities we had stated in chapter 4, what has made the aforementioned fibre cavity device implementations possible? By summarising the advances in the field, we can see how. Let us do so:⁷

- *Birefringence*: fibre-tip mirrors with negligible birefringence were created by using several shorter ablation pulses and rotating the fibre between each pulse, as opposed to one longer pulse [135]. The same method can be used to deliberately engineer a desired birefringence into the cavity mode [202], which allows us to make even more efficient devices that utilise Purcell enhancement [102] and already was used to facilitate a combined ‘qubit’ and ‘herald’ crossed cavity pair [110].
- *Mirror curvature*: minimising power fluctuations in the CO₂ ablation laser greatly improved homogeneity of fibre-tips and accuracy and precision of the desired curvature ζ [132, 203, 204]. FIB milling has also been used to minimise ζ to just 5 μm [122], giving smaller mode volumes and therefore stronger atom-photon coupling (although doing so makes the maximum cavity length shorter too, reducing optical access).
- *Fibre loss*: using photonic crystal fibres or splicing gradient-index lenses onto the fibre-tips has allowed for much improved mode-matching [205, 206].⁸ Moreover, this meant the cavity length could increase to the order of millimetres [132]; these longer cavities, with their good optical access, have in turn improved atom and ion trapping implementations [66, 113].
- *Vacuum compatibility*: Historically, researchers noticed a degradation of fibre-tip mirrors when in vacuum. The longer the cavity was left in, the more the finesse would degrade [109, 174]. Recently however, this was identified as a chemically-induced loss,

⁷ Birefringence is defined in section 4.3.2; mirror curvatures and fibre loss in section 4.1; and frequency locking in section 5.2.4.

⁸ With respect to photonic crystal fibres, they are incompatible with the ablation itself [134]: instead, they are spliced to regular fibres.

which was then rectified in the vacuum assembly process, and so fibre cavities are now fully vacuum-compatible in the long term [207].

- *Frequency locking*: following the first successful lock using the PDH method [129], more recent implementations have now been successfully locked (in spite of the translation stages that make the cavity less stable in vacuum [208, 209]).

With respect to our own cavity, we know from these advances that frequency locking is intrinsically possible—we mentioned already in chapter 5 that replacing the cavity mounting finger should allow us to do so. What this and the other benefits enable is reliability and repeatability. We have already assembled an SM-SM fibre cavity with submicron precision in vacuum. Moreover, we achieved this with a batch of fibre-tips that exhibited more manufacturing variations than those mentioned above (see table 4.1). In this context, we see that today, buoyed by our and others' advances we have the ability to assemble many of our fibre cavities in a reproducible way: our assembly technique is the only current one that can be applied to assemble multiple cavities on the same (V-groove) chip. Even the latest techniques for making monolithic chips containing fibre cavities, namely using wire cutting [141] or by 3D printing glass [210], are not capable of doing so with submicron precision. Reliability is certainly a key requirement for novel quantum devices; by investigating the issues and solving them as we had done extensively in chapter 4, we have therefore made steps towards realising reliable fibre cavity devices in the future. For example, one such device could be our envisaged 2D tweezer implementation.

Connecting fibre cavities

To bring together all of the above discussions, we display in figure 6.1 a speculative idea of an experiment our cavity could enable. By connecting many cavities with an optical




Figure 6.1: Possible network topologies. The $N \times N$ switch connects N cavities to N other cavities. From this, we can investigate whether directional or symmetric cavities are more efficient in networking, and in particular whether the quantum optical description of dissipation from a driven cavity matches the classical description [105].

switch as realised by Kim et al. [177, 211], we can investigate the nature of coupling cavities past the elementary size. In this context, an optical switch can be thought of as a railway switchyard for light, where light on any line (input) can be directed to any other line (output) at will. Although, due to the losses involved, making a network of many cavities may be pushing against formidable limits in experimental resources, it is still achievable today: our simulations and experiments have indicated what direction we should follow to pursue this.

In particular, we can investigate the connection of multiple symmetric fibre cavities. The $N \times N$ switch is so called because it connects N cavities to N other cavities.⁹ We stress that what figure 6.1 depicts is only illustrative: we do not have to use N cavities, we may use only a few. What we can do is use this network to examine the deterministic nature of photon emission and absorption when an atom is coupled to a cavity.

This deterministic aspect is one we had mentioned in chapter 4: we constructed a fully deterministic atom-photon interface. When concerning atom-photon interfacing, what is ‘fully deterministic’? Usually, we take this to mean the case where an atom is trapped for

⁹ It functions by using an SLM and two lenses: the first lens focuses the light from the N inputs onto the SLM, Fourier transforming it [42]. On the opposite side, another lens at the same distance away and focused on the SLM performs the inverse Fourier transform. Altogether, this transforms the SLM’s phase or amplitude modulation to a spatial modulation—allowing us to redirect light from one of the N inputs to any of the N outputs.

much longer than interfacing time, so that we may interface many times at our will before we lose the atom. However, this definition does not account for decoherence. Crucially, for a directional cavity we have no ‘built-in’ way of checking whether the emitted photon (or the photon to be absorbed) is incoherent until we complete the network operation. Recall from chapter 1 our practical definition: decoherence means the operation happens, but with an undesired outcome. For a directional cavity, this means if a small amount of light leaks through the high-reflection mirror, decoherence on that side/linking channel of the cavity may possibly affect the quantum state of any light leaking through the high-transmission mirror. But are these two states correlated in any way, or not? It is difficult to determine this with a directional cavity, because it is difficult to measure the state of light leaking through the high-reflection mirror.

For this reason, we need a practical way of declaring the outcome invalid, so that we may try again for a valid outcome. Luckily, for networks of fibre cavities, we have a recent breakthrough with respect to this. We may build a crossed-cavity node—two fibre cavities crossed at right angles to each other—and use one cavity as a ‘qubit’ cavity for interfacing with photonic qubits, and the other ‘herald’ cavity for non-destructive detection of these photonic qubits [29, 110].

The devices in [29, 110] use directional cavities. But, using our symmetric cavity design, we can postulate another simpler approach. We dedicate one side of the cavity as the ‘herald’ side, and the other side as the ‘qubit’ side. We then connect the herald sides of two cavities directly, as shown in figure 6.1. Because the cavity is symmetric, we can then use the techniques above to determine whether decoherence on one channel affects the state on the other channel. In this way, we can fully determine the nature of emission and coupling of a two-sided symmetric cavity. This is what we mean by ‘fully deterministic’: we can test whether interfacing between atoms and photons is affected by its environment, while

subtracting the effect of loss to the environment. With respect to examining the decoherence, first we would need to prepare a highly quantum state, such as a superposed or entangled state. If at least one of the cavity modes uses our 2D atom array design, we could in the future potentially entangle many atoms at once by utilising collective quantum jumps of well-coupled atoms [212], or by moving atoms in and out of the cavity [213].

We close with a comment on some very recent research in our group. Our group has performed quantum feedback on the state of a photon emitted from the cavity. Because the photon temporal lengths are long, we can perform feedback on a second photon using our detection result from the first photon. By considering this time-resolved interference at a beam splitter, we could ensure that the second photon always exits through the same beam splitter output [214]. This resembles a ‘Schrödinger cat’ situation: the second photon, the ‘cat’, is both alive and dead until we measure it. But, because the feedback correlates its state with the first photon’s state, our feedback effectively provides the antidote to keep the cat alive. In our laboratory, progress is currently under way to connect two distant cavities with a fibre link. If we therefore consider our symmetric two-sided fibre cavity, or a similar one, in the same fashion, we may therefore examine the quantum nature of deterministic cavity emission at a more fundamental level than by examining the output of a single photon that may be emitted from both sides of a symmetric optical cavity [105].

Conclusion

The overall goal of this thesis was to design, implement and optimise a [fibre Fabry-Pérot cavity \(FFPC\)](#) for use as an efficient, faithful atom-photon interface. Motivated by its suitability for quantum networks of atoms and photons, we successfully constructed a device with three key features for this purpose: a small mode volume for strong atom-photon coupling, directly fibre-coupled cavity light, and optical access for trapped atoms.

We began by optimising the parameters of a single atom-cavity system for this purpose, and then optimised two connected together—an elementary cavity-based quantum network. By simulating the single photon generation/absorption [V-STIRAP](#)¹ process then, using a master equation approach, we numerically optimised the system parameters using the Nelder-Mead algorithm and while keeping the [STIRAP](#)-driving laser pulse shape fixed. We used the algorithm with a performance function that gives us a figure of merit for optimal networking processes. At this point, we obtained a 95.4 % photon generation efficiency for realistic parameters for an [FFPC](#).

After finding the optimal system parameters, we then optimised the laser pulse shape. We firstly used the Krotov algorithm, obtaining a suboptimal final efficiency and a spike in the laser shape at the beginning of the pulse. From this we deduced that, because the algorithm

¹ See section [2.2.3](#) for a description of what [STIRAP](#) is.

tried to minimise the cavity decay even though [V-STIRAP](#) relies on it, Krotov and other gradient-based algorithms are unsuitable for optimising cavity-based quantum networking processes that rely on decay. In our second attempt, we therefore applied the Nelder-Mead method to optimise pulse points through which we fit splines. This was successful: we obtained a 99.0 % generation efficiency with a process coherence of 99.8 %. Further, we modified the algorithm to find the most efficient shapes to produce photons of given temporal lengths, able to keep the process efficient right down to the adiabatic limit.

Finally, we optimised pulse shapes for remote state transfer: a photon generation followed by absorption process in two connected cavities. Simulating two identical copies of our optimal cavity with similar parameters to the fibre cavity we constructed, we achieved a 39 % transfer efficiency. We found that, due to the fixed coupling rates of impinging photons to the cavity, we must use pulse shapes that sacrifice generation efficiency for absorption efficiency to keep the overall process efficient—when we modelled this, the transfer efficiency dropped to 20 %. We also successfully applied our technique to a hybrid network with different atom-cavity parameters, thus modelling the realistic scenarios in the laboratory wherein manufacturing differences mean every network node is different. Using the parameters of the macroscopic cavity in our other laboratory, we achieved 11 % efficiency, mostly limited by the atom in the macroscopic cavity re-emitting the impinging photon, which then leaked back out.

We then progressed with our experimental work. We constructed the first [FFPC](#) formed of two [single-mode \(SM\)](#) fibres in vacuum. At $(90 \pm 2) \mu\text{m}$ long and with fibre-tip diameters of $125 \mu\text{m}$, our cavity has the optical access for trapped atoms in optical tweezers, allowing for deterministic coupling of atoms to the cavity mode. By applying mode-matching theory specific to [FFPCs](#), we designed a cavity with the optimal mode-matching efficiency of 64 %

for the best fibre-tips in our batch.² However, due to breakage of the fibres, we instead used a pair where one fibre-tip's concave mirror feature was decentred by $2.5\ \mu\text{m}$, affecting the mode-matching. After assembly, we measured a 34 % mode-matching efficiency, and confirmed the decentration by temporarily misaligning the cavity and measuring the resulting asymmetric cavity lineshape. This compares well with other cavities formed of an [SM](#) and a [multi-mode \(MM\)](#) fibre, which typically exhibit mode-matching efficiencies of around 60 % on the input only [129], and thus $0.6^2 = 0.36$ total mode-matching efficiency.

We then assembled the cavity in vacuum. Here, our major advance was to devise an assembly technique that produces a compact, permanently aligned microcavity with sub-micron precision. Two features made the assembly technique perform well. Firstly, we minimised epoxy-induced misalignment by testing different epoxy types and choosing the optimum epoxy type based on its parameters and the results of simulated vacuum bakeout tests. Secondly, we used an assembly procedure that works with rather than against gravity and allowed for a suitably thin spacer to be fit within the assemblage, thus mitigating misalignment caused by both gravity and finite manufacturing tolerance. Our assembly technique will therefore be of benefit to other microassemblies in vacuum. Overall, from the assembled cavity parameters we predicted an atom-cavity co-operativity of 94 ± 5 : well into the strong coupling regime, nearly two orders of magnitude better than other macroscopic cavities created in our research group [55] and double that of the best performing fibre cavities of similar type [174].³

After assembling the cavity, we implemented the cavity in our atom-photon coupling experiments. We designed this experiment to facilitate multiple trapped and movable single atoms in the cavity. In this design, holographically-generated optical tweezers from a [spatial](#)

² The fibre-tips were manufactured in a collaboration between the Reichel group in ENS Paris and a member of our group before the author's time.

³ See table 4.3 for a full summary of the cavity parameters.

[light modulator \(SLM\)](#) give us a 2D reconfigurable array of atoms in the cavity mode. Our focus was on implementing the cavity rather than the tweezers; we implemented a probabilistic loading scheme from which atoms are either pushed from a [magneto-optical trap \(MOT\)](#), or we drop the [MOT](#) through the cavity by gravity. We then looked for Purcell-enhanced fluorescence from the cavity. Because our measurements were hindered by mechanical instability that we determined to be caused by the cavity holding finger within the vacuum assembly, we were not able to frequency-lock the cavity in the usual way when performing these experiments. Instead, we stabilised the frequency using a lock-in amplifier while scanning about the cavity resonance. In spite of this, we were successful in detecting Purcell-enhanced fluorescence from a few atoms, an effect that can only be due to atom-cavity coupling. By replacing the mechanically unstable mounting finger, we therefore expect that this cavity would be capable of atom-photon interfacing well into the strong coupling regime.

Overall, the work in this thesis facilitates deterministically controlled, strongly-coupled and networked quantum systems of atoms in cavities. We expect that the findings and techniques presented here will be particularly useful in guiding the next generation of cavity-based atom-photon interfacing experiments and quantum networks, providing a platform for new, novel experiments in quantum physics at the fundamental level.

Final comments

Here, the author gives a more personal view over the work in this thesis.

All around the world, new research has created devices that push our understanding of the quantum world to ever higher levels, giving us a toolbox with which we can probe our understanding at the most fundamental level [39]. With curiosity and perseverance, we have built devices such as the frequency comb [215] and gravitational wave interferometer [216]. Remarkably, the latter is built using giant optical cavities. In fact, there is a direct connection

between the gravitational wave interferometer, the CQED research in this thesis, and the concept of a quantum internet: Kimble had a role in all of them, both by his own research and through collaboration [14, 54, 217].

Physics today is a global, collaborative effort. Although the device presented in this thesis was constructed solely by the author, the promise that it shows for investigating atom-photon interactions well into the strong coupling regime, together with the current exciting research in the CQED field [21], suggests that if we keep improving with regards to sharing our research and fostering collaboration and ideas, we will collectively make new, fundamental discoveries that benefit all of us.

How to calculate the mixing angles of two- and three-level systems

When an atom and light interact, the energy levels of the coupled system may shift, as discussed in chapter 2. If the frequency of the light involved (or the emitter frequency) changes, then the energy shifts will change too. We can use mixing angles to concisely define this frequency dependence of the energies and the system's (dressed) states. In this appendix, we show the full workings of how to calculate the mixing angles of the two-level and three-level atom-cavity systems in chapter 2. The basic idea is to pull out constant factors from the system eigenvectors until we can write them in a simpler way.

A.1 Two-level system

Starting from the energy eigenfrequencies and eigenstates for the off-resonantly driven two-level system (equations 2.44 and 2.45), we firstly simplify them by defining

$$\omega_{\Delta} = \pm \frac{1}{2} \sqrt{\Delta^2 + 4g^2} = \pm \frac{1}{2} \sqrt{\Delta^2 + \Omega^2} = \pm \frac{\tilde{\Omega}}{2}, \quad (\text{A.1})$$

where

$$\tilde{\Omega} = \sqrt{\Delta^2 + 4g^2} \quad (\text{A.2})$$

for a Rabi frequency $\Omega = 2g$. Substituting this into the eigenvectors gives

$$|+, 1\rangle = \frac{\Delta + \tilde{\Omega}}{-\Omega} |g, 1\rangle + |e, 0\rangle, \quad (\text{A.3})$$

$$\begin{aligned} |-, 1\rangle &= \frac{\Delta - \tilde{\Omega}}{-\Omega} |g, 1\rangle + |e, 0\rangle \\ &= \frac{(\Delta - \tilde{\Omega})(\Delta + \tilde{\Omega})}{-\Omega(\Delta + \tilde{\Omega})} |g, 1\rangle + |e, 0\rangle \\ &= \frac{\Delta^2 - \tilde{\Omega}^2}{-\Omega(\Delta + \tilde{\Omega})} |g, 1\rangle + |e, 0\rangle \\ &= \frac{-\Omega^2}{-\Omega(\Delta + \tilde{\Omega})} |g, 1\rangle + |e, 0\rangle \\ &= \frac{\Omega}{\Delta + \tilde{\Omega}} |g, 1\rangle + |e, 0\rangle. \end{aligned} \quad (\text{A.4})$$

If we pull out a constant factor from each eigenvector:

$$|+, 1\rangle = \frac{1}{-\Omega} ((\Delta + \tilde{\Omega})|g, 1\rangle - \Omega|e, 0\rangle), \quad (\text{A.5})$$

$$|-, 1\rangle = \frac{1}{\Delta + \tilde{\Omega}} (\Omega|g, 1\rangle + (\Delta + \tilde{\Omega})|e, 0\rangle), \quad (\text{A.6})$$

then upon inspection, we can define

$$\tan \theta = \frac{-\Omega}{\Delta + \tilde{\Omega}} = \frac{o}{a}, \quad \sin \theta = \frac{o}{h}, \quad \cos \theta = \frac{a}{h}, \quad (\text{A.7})$$

and set h to whatever factor we pulled out of equation A.5 or equation A.6¹ to get equations 2.46 and 2.47 in the main text:

$$|+, 1\rangle = \cos \theta |g, 1\rangle + \sin \theta |e, 0\rangle, \quad (\text{A.8})$$

$$|-, 1\rangle = -\sin \theta |g, 1\rangle + \cos \theta |e, 0\rangle.$$

$$\tan \theta = \frac{-2g}{\Delta + \sqrt{\Delta^2 + 4g^2}}. \quad (\text{A.9})$$

A.2 Three-level Lambda-system

Starting from equation 2.61 for the three-level Λ -system, and on Raman resonance $\Delta = \Delta_L = \Delta_C$ we calculate the three energy eigenfrequencies:

$$\lambda_0 = -\Delta, \quad (\text{A.10})$$

$$\lambda_{\pm} = \frac{-\Delta \pm \sqrt{4g^2 + \Delta^2 + \Omega^2}}{2}, \quad (\text{A.11})$$

and the three eigenvectors:

$$|\lambda_0\rangle = \frac{-2g}{\Omega} |u, 0\rangle + |g, 1\rangle, \quad (\text{A.12})$$

$$|\lambda_{\pm}\rangle = \frac{\Omega}{2g} |u, 0\rangle + \frac{\Delta \mp \sqrt{4g^2 + \Delta^2 + \Omega^2}}{2g} |e, 0\rangle + |g, 1\rangle. \quad (\text{A.13})$$

To simplify them, we set

$$\tilde{\Omega}_{\Lambda} = \sqrt{4g^2 + \Delta^2 + \Omega^2} = \sqrt{\Omega_C^2 + \Delta^2 + \Omega_L^2} \quad (\text{A.14})$$

¹ We are free to set h as we like, not only $\sqrt{(\Delta + \tilde{\Omega})^2 + (-\Omega)^2}$, since no matter what factor we pull out, the state will still be an eigenvector.

where $\Omega_C = 2g$ is the cavity coupling Rabi frequency and the laser coupling now has an ‘L’ subscript so we can distinguish the two couplings. Substituting into the eigenvectors gives

$$|\lambda_0\rangle = -\frac{\Omega_C}{\Omega_L}|u, 0\rangle + |g, 1\rangle, \quad (\text{A.15})$$

$$|\lambda_{\pm}\rangle = \frac{\Omega_L}{\Omega_C}|u, 0\rangle + \frac{\Delta \mp \tilde{\Omega}_\Lambda}{\Omega_C}|e, 0\rangle + |g, 1\rangle. \quad (\text{A.16})$$

Firstly, we deal with $|\lambda_0\rangle$. Pulling out a factor,

$$|\lambda_0\rangle = \frac{1}{-\Omega_L} (\Omega_C|u, 0\rangle - \Omega_L|g, 1\rangle), \quad (\text{A.17})$$

then upon inspection, we can define

$$\tan \Theta = \frac{-\Omega_L}{\Omega_C} = \frac{o}{a}, \quad \sin \Theta = \frac{o}{h}, \quad \cos \Theta = \frac{a}{h}, \quad (\text{A.18})$$

and normalise equation A.17 using

$$h = \sqrt{\Omega_L^2 + \Omega_C^2} \quad (\text{A.19})$$

to get:

$$|\lambda_0\rangle = \cos \Theta|u, 0\rangle + \sin \Theta|g, 1\rangle. \quad (\text{A.20})$$

Secondly, we deal with the remaining two eigenvectors, $|\lambda_{\pm}\rangle$. When we used equation A.19, we had set h as a normalisation constant when defining the Θ mixing angle. We can pull this factor out of $|\lambda_{\pm}\rangle$ and discard it, and $|\lambda_{\pm}\rangle$ will remain an eigenvector. Then, substituting the expressions for Θ in, we get

$$|\lambda_{\pm}\rangle = \frac{1}{\Omega_C} (\sin \Theta|u, 0\rangle + (\Delta \mp \tilde{\Omega}_\Lambda)|e, 0\rangle - \cos \Theta|g, 1\rangle). \quad (\text{A.21})$$

We can further simplify $|\lambda_{\pm}\rangle$ as follows, using equation A.14:

$$\begin{aligned}
 \Delta - \tilde{\Omega}_{\Lambda} &= \frac{(\Delta - \tilde{\Omega}_{\Lambda})(\Delta + \tilde{\Omega}_{\Lambda})}{\Delta + \tilde{\Omega}_{\Lambda}} \\
 &= \frac{\Delta^2 - \tilde{\Omega}_{\Lambda}^2}{\Delta + \tilde{\Omega}_{\Lambda}} \\
 &= \frac{-(\Omega_L^2 + \Omega_C^2)}{\Delta + \tilde{\Omega}_{\Lambda}} \\
 &= \frac{-h^2}{\Delta + \tilde{\Omega}_{\Lambda}}.
 \end{aligned} \tag{A.22}$$

Using this relation, equation A.21 becomes

$$|\lambda_{+}\rangle = \frac{1}{\tilde{\Omega}_{\Lambda} + \Delta} \left((\tilde{\Omega}_{\Lambda} + \Delta) \sin \Theta |u, 0\rangle - h^2 |e, 0\rangle - (\tilde{\Omega}_{\Lambda} + \Delta) \cos \Theta |g, 1\rangle \right), \tag{A.23}$$

$$|\lambda_{-}\rangle = \frac{1}{-h^2} \left(-h^2 \sin \Theta |u, 0\rangle - (\tilde{\Omega}_{\Lambda} + \Delta) |e, 0\rangle + h^2 \cos \Theta |g, 1\rangle \right). \tag{A.24}$$

Upon inspection, if we define

$$\tan \Phi = \frac{-h^2}{\tilde{\Omega}_{\Lambda} + \Delta} = \frac{o'}{a'}, \quad \sin \Phi = \frac{o'}{h'}, \quad \cos \Phi = \frac{a'}{h'}, \tag{A.25}$$

then we can pull out another constant h' . Afterwards, we get equations 2.62–2.64 in the main text. These are, respectively:

$$\begin{aligned}
 |\psi^0\rangle &= \cos \Theta |u, 0\rangle + \sin \Theta |g, 1\rangle, \\
 |\psi^+\rangle &= \sin \Theta \cos \Phi |u, 0\rangle + \sin \Phi |e, 0\rangle - \cos \Theta \cos \Phi |g, 1\rangle, \\
 |\psi^-\rangle &= \sin \Theta \sin \Phi |u, 0\rangle - \cos \Phi |e, 0\rangle - \cos \Theta \sin \Phi |g, 1\rangle,
 \end{aligned} \tag{A.26}$$

then again substituting equation A.14,

$$\tan \Theta = \frac{\Omega}{-2g}, \quad \tan \Phi = \frac{-\sqrt{4g^2 + \Omega^2}}{\sqrt{4g^2 + \Omega^2 + \Delta^2} + \Delta}, \quad (\text{A.27})$$

and finally, adding in the cavity energies to get the energy eigenfrequencies (for n quanta),

$$\begin{aligned} \omega_0 &= \omega_C \left(n + \frac{1}{2} \right), \\ \omega_{\pm} &= \omega_C \left(n + \frac{1}{2} \right) + \frac{\Delta \mp \sqrt{4g^2 n + \Omega^2 + \Delta^2}}{2}. \end{aligned} \quad (\text{A.28})$$

Liouville space

When using Krotov's algorithm in section 3.4, we need to run a time-reversed evolution of our atom-cavity system for the optimising trajectory. However, our system is dissipative. It has cavity decay through the mirrors, and spontaneous decay of the atom to the environment. We cannot reverse a dissipation. Therefore, we move from Hilbert space to Liouville space, which allows us to include decays in Schrödinger's equation and use that to calculate the evolution. In this appendix, we briefly demonstrate what Liouville space is.

Let us consider a two level atom with ground state $|g\rangle$ and excited state $|e\rangle$. Starting from the master equation (3.1), consider one path of decay: spontaneous emission by the atom, $\sqrt{2\gamma}|g\rangle\langle e|$. In the basis $\{|g\rangle, |e\rangle\}$, the corresponding collapse operator is

$$\hat{C}_1 = \sqrt{2\gamma} \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}. \quad (\text{B.1})$$

Using equation 3.2 gives the corresponding dissipation superoperator as

$$\hat{D}_1(\hat{\rho}) = \gamma \begin{pmatrix} 2\rho_{22} & -\rho_{12} \\ -\rho_{21} & -2\rho_{22} \end{pmatrix}, \quad (\text{B.2})$$

where ρ_{ij} is the element of the density matrix $\hat{\rho}$ in its i^{th} row and j^{th} column. Because \hat{D} is non-unitary,¹ an interaction between the system and environment irreversibly changes the mixture of the states that comprise $\hat{\rho}$. Therefore, it would not be possible to evolve $\hat{\rho}$ without keeping track of how this mixture is changed.

One way to circumvent this limitation is to use the Lindblad master equation 3.1, as we do for all of chapter 3 except section 3.4 on Krotov. But, this does not allow for time-reversed propagation, as we do not know from which path the decay came from. So instead, we switch to Liouville space as follows. Firstly, we input \hat{D}_1 from equation B.2 into the master equation 3.1 to get the equation of motion in matrix form,

$$\begin{pmatrix} \dot{\rho}_{11} & \dot{\rho}_{12} \\ \dot{\rho}_{21} & \dot{\rho}_{22} \end{pmatrix} = \gamma \begin{pmatrix} 2\rho_{22} & -\rho_{12} \\ -\rho_{21} & -2\rho_{22} \end{pmatrix}, \quad (\text{B.3})$$

where $\dot{\rho}_{ij} = \frac{d\rho_{ij}}{dt}$. Secondly, we write each $\dot{\rho}_{ij}$ in terms of the the matrix elements on the right-hand side,

$$\begin{aligned} \dot{\rho}_{11} &= \gamma (2\rho_{22}), \\ \dot{\rho}_{12} &= \gamma (-\rho_{12}), \\ \dot{\rho}_{21} &= \gamma (-\rho_{21}), \\ \dot{\rho}_{22} &= \gamma (-2\rho_{22}). \end{aligned}$$

¹ This means $\hat{D} \neq \hat{D}^\dagger$.

Thirdly, we write these coupled equations in terms of ρ_{ij} , giving

$$\begin{pmatrix} \dot{\rho}_{11} \\ \dot{\rho}_{12} \\ \dot{\rho}_{21} \\ \dot{\rho}_{22} \end{pmatrix} = 2\gamma \begin{pmatrix} 0 & 0 & 0 & 2 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -2 \end{pmatrix} \begin{pmatrix} \rho_{11} \\ \rho_{12} \\ \rho_{21} \\ \rho_{22} \end{pmatrix}. \quad (\text{B.4})$$

This is now in Liouville space. The matrix $\hat{\rho}$ has been ‘folded out’ onto a vector, and the dissipation superoperator \hat{D}_1 has been written in a larger form of $\dim(N \times N) = 4$, where N is the size of our system.²

For our more complex atom-cavity system, we may apply the same principle as above to all the possible decay paths. Then we may combine the equation of motion and the dissipation superoperator in equation 3.1 to write it in Liouville space as³

$$\begin{aligned} \frac{d|\rho(t)\rangle\rangle}{dt} &= (-i\hat{L} + \hat{D}) |\rho(t)\rangle\rangle = -i(\hat{L} - i\hat{D}) |\rho(t)\rangle\rangle \\ &= \hat{\mathcal{L}} |\rho(t)\rangle\rangle, \end{aligned} \quad (\text{B.5})$$

where

$$\hat{\mathcal{L}} = -i [\hat{H}, |\hat{\rho}(t)\rangle\rangle], \quad (\text{B.6})$$

and where $|\rho(t)\rangle\rangle$ is the vector on the right-hand side of equation B.4 and is double-braced to highlight that it is in Liouville space. We have omitted the operators’ dependencies on $|\rho(t)\rangle\rangle$ to simplify the notation. Note that this form is the same as Schrödinger’s equation, but with H replaced by $i\hat{\mathcal{L}}$.

² Note also the state vector is of length $N \times N = 4$.

³ We set $\hbar = 1$, like in chapter 3.

To perform a normal time evolution, we solve equation B.5 to get

$$\hat{U}(T) = e^{\int_0^T \hat{\mathcal{L}}(t) dt}, \quad (\text{B.7})$$

cf. equation 2.31 and Sakurai [51]. How do we execute a time-reversed evolution? Just like with Schrödinger's equation, we take the Hermitian conjugate, but also flip the integral limits, yielding

$$\hat{U}(T)^\dagger = e^{-\int_T^0 \hat{\mathcal{L}}(t)^\dagger dt}. \quad (\text{B.8})$$

We can do this because \hat{L} is Hermitian and \hat{D} is anti-Hermitian, $\hat{D}^\dagger = -\hat{D}$, as proven in [83]. Therefore, although we evolve back from time t to 0, the decays act in exactly the same way as if we were going from time 0 to t .

For our purpose, Liouville space effectively works the same as Hilbert space, but allows us to group the collapse operators and the equation of motion for the density matrix (von Neumann equation B.6) into one convenient term. Then, we can make a time-reversed evolution with decays.

We stress that when including decays, backward propagation no longer corresponds to backward evolution of the state. A state that decays while propagating forward will not revert the decay further while propagating backward. The system will decay in the reverse-time direction just like in forward-time, so after one forward trajectory and one optimising trajectory, the system will decay 'twice'.

Theory of a 2D array of single-atom optical tweezers

In section 5.2.1, we set out our design for combining the fibre cavity with a 2D array of single atoms. One optical tweezer traps a single atom; we can generate these tweezers holographically to get the 2D array. By updating the hologram to reconfigure the array, each atom may be moved independently of one another. To better understand the constraints of the tweezers and get the best performance from them, we need to understand their principle of operation. So firstly, we briefly explain the theory of how we can trap the atoms. Afterwards, we explain how we can generate the 2D array of single atoms in the cavity.

C.1 Theory of single-atom optical tweezers

When we refer to optical tweezers, we refer to tightly-focused laser beams that trap particles in a potential well created by the gradient of the beam intensity. It is well known that the gradient of a conservative potential U results in a force

$$\mathbf{F} = -\nabla U. \tag{C.1}$$

Now, we want to calculate this force. One may think of the famous experiments by Ashkin and Dziedzic at the microscale where biological cells [218] and glass spheres [219] were levitated and manipulated with optical tweezers. For these cases where the object is much larger than the beam wavelength, the origin of the trapping force comes from a radiation pressure conservation argument [220, 221]. For single atoms, in fact the origin of the trapping force is the dipole interaction with the laser. Remarkably, we can use the theory we already derived for atom-cavity interactions in section 2.2 to immediately calculate the resulting trapping force, known as the dipole force.

Our physical picture is a two-level atom interacting with a laser in free space, rather than a cavity field. Starting from equation 2.43, we simply replace the atom-cavity coupling with an atom-laser coupling to get¹

$$\hat{H}_d = -\frac{\hbar\Delta_L\hat{\sigma}_z}{2} + \frac{\hbar\Omega_L}{2}(\hat{\sigma}^+ + \hat{\sigma}^-). \quad (\text{C.2})$$

where we have replaced Δ with the laser detuning Δ_L and g with the laser Rabi frequency $\Omega_L/2$. We are working in a semiclassical picture—no need to quantise the light field here—hence, we have dropped the photon operators \hat{c} and \hat{c}^\dagger and the photon numbers in the eigenstates $|+, 1\rangle$ and $|-, 1\rangle$ (equation 2.46): the eigenstates are now just $|+\rangle$ and $|-\rangle$. Otherwise, this is exactly the same system as before: even figure 2.6 is similar. The $|+\rangle$ eigenstate has eigenenergy

$$\frac{E_d}{\hbar} = \frac{1}{2}\sqrt{\Delta_L^2 + \Omega_L^2}. \quad (\text{C.3})$$

The key property of optical tweezers for atoms is that the beam is far detuned.² Consider a far

¹ This is in the interaction picture, cf. equation 2.58 in the Schrödinger picture.

² Hence, they are sometimes referred to as far-off-resonant traps (FORTs).

off-resonant laser, $\Delta_L \gg \Omega_L$. In this case, we can Taylor-expand the square root to first order:

$$\frac{E_d}{\hbar} \simeq \frac{\Delta_L}{2} \left(1 + \frac{\Omega_L^2}{2\Delta_L^2} \right), \quad (\text{C.4})$$

and in tandem get (see equation 2.47)

$$|+\rangle \simeq |g\rangle,$$

and so if we consider just the energy shift due to the laser—the second term in equation C.4—we get³

$$U_d = \frac{\hbar\Omega_L^2}{4\Delta_L} \quad (\text{C.5})$$

for $|g\rangle$. For $|e\rangle$, we can repeat the calculation for the $|-\rangle$ eigenstate to find that U_d is equal in magnitude but with the opposite sign.⁴ In these two cases, the energy level shift is known as the linear Stark shift [147].⁵ We simply input equation C.5 to equation C.1 to get our trapping force.

From $\mathbf{F}_d = -\nabla U_d$, we can see why we need the tweezers to be tightly focused: because Ω_L is proportional to the beam's intensity, then the steeper the spatial variation in intensity, the larger \mathbf{F}_d is. Furthermore, we must also make sure that $\Delta_L < 0$ (red-detuning), so that \mathbf{F}_d is positive and pushes the atom towards the focus of the beam—its most intense point, and the bottom of the potential well.

³ From footnote 24 in chapter 2, we might expect the RWA to be invalid because of the large detuning. However, considering that we keep the atom trapped for a long time, the timescale of the interaction is long and we quickly average the interaction over enough cycles for the RWA to be valid again [222].

⁴ The switch in sign when in state $|e\rangle \simeq |-\rangle$ implies that the tweezers would not work for an excited atom. But, for our intended experiments, we will generally stay well below the saturation limit, because we will use weak or single-photon laser pulses. Thus, the trapped atom is almost never in state $|e\rangle$.

⁵ Sometimes referred to as AC Stark shift.

How much force do we need to trap an atom? It is easier to calculate instead the trap depth,

$$T_d = \frac{U_d}{k_B}, \quad (\text{C.6})$$

where k_B is Boltzmann's constant. We see that T_d is simply how deep the potential well is, expressed as a temperature; if the atom's temperature exceeds T_d , it will escape the trap. Now, we can further simplify our question to: what U_d do we need to trap an atom at temperature T ? To calculate U_d , we need an expression for Ω_L , which is [42]

$$\Omega_L^2 = \frac{\Gamma^2 I}{2I_{\text{sat}}}, \quad (\text{C.7})$$

where $\Gamma = 2\gamma$ is the atomic linewidth, I is the tweezer beam intensity, and I_{sat} is the atomic saturation intensity. So now, we need expressions for I and I_{sat} , which are at the focus of the beam with waist w_0 [42]

$$I = \frac{2P}{\pi w_0^2}, \quad (\text{C.8})$$

$$I_{\text{sat}} = \frac{\hbar\omega_a^3\Gamma}{12\pi c_0^2}, \quad (\text{C.9})$$

where P is the beam power,⁶ ω_a is the atomic transition frequency and c_0 is the speed of light. Substituting equations C.5, C.7, C.8 and C.9 into equation C.6 yields our desired result,

$$T_d = \frac{3c_0^2\Gamma P}{w_0^2\omega_a^3\Delta_L}. \quad (\text{C.10})$$

⁶ The factor of 2 in equation C.8 arises from the fact that I of a Gaussian beam scales as e^{-2r^2/w_0^2} from the waist. If we divide that by the area $A = \pi r^2$ and use l'Hôpital's rule, a factor of 2 drops out. Then using $P = IA$, we get equation C.8.

C.2 Generating a reconfigurable 2D array of single atoms

We now explain how we generate a 2D array of single atoms in the cavity with our experimental setup.

Firstly, we need to check that a single tweezer is compatible with the cavity. What is T_d if we use a $\text{NA} = 0.60$ lens, and tweezers with $\lambda = 1064 \text{ nm}$ and $P = 100 \text{ mW}$? Using equation C.10, we get $T_d = 47.4 \text{ mK}$.⁷ In this regime, it is therefore highly unlikely to trap atoms at room temperature, because they will easily escape the potential well. Therefore, this is why we need to laser-cool the atoms before loading them into the tweezers.

After doing so, how do we ensure that we only have a single atom in a tweezer? This we achieve ‘for free’ via the collisional blockade effect [223, 224]. Put simply, if there are initially multiple atoms in a tweezer, then pairs of atoms will collide and be lost from the trap until eventually only one or zero atoms remain. We operate in this regime because our w_0 is small enough ($\lesssim 1 \mu\text{m}$) and therefore light-assisted collisions of atom pairs are the dominant loss process [225]. Assuming the collisional blockade effect dominates, we would expect half of the traps to have an atom on average. We can increase this filling fraction further by applying feedback to the trap depth depending on if an atom is present/detected [49], or by reassembling the array of tweezers one-by-one until all the desired sites are filled [226, 227].

We had mentioned ‘array’ just now. Expanding on this, we now describe how we form a reconfigurable 2D array of trapped atoms. Recall that in chapter 4, we discussed how we can move atoms in and out of the cavity mode using an SLM to holographically generate a 2D array of tweezers [107, 108]. Figure C.1 illustrates the experimental setup for this. In our design, our SLM⁸ consists of a 2D array of phase-modulating micro-pixels. A 20 W

⁷ If we consider not just a two-level atom but the multi-level structure of the D_2 line, we correct this by a factor 1.42 and have $T_d = 33.4 \text{ mK}$. For brevity, we have restricted our analysis to a two-level atom, but further details of the correction can be found in our own group’s research [49, 153] or by Neuzner et al [53].

⁸ Meadowlark Optics XY Series P512-106

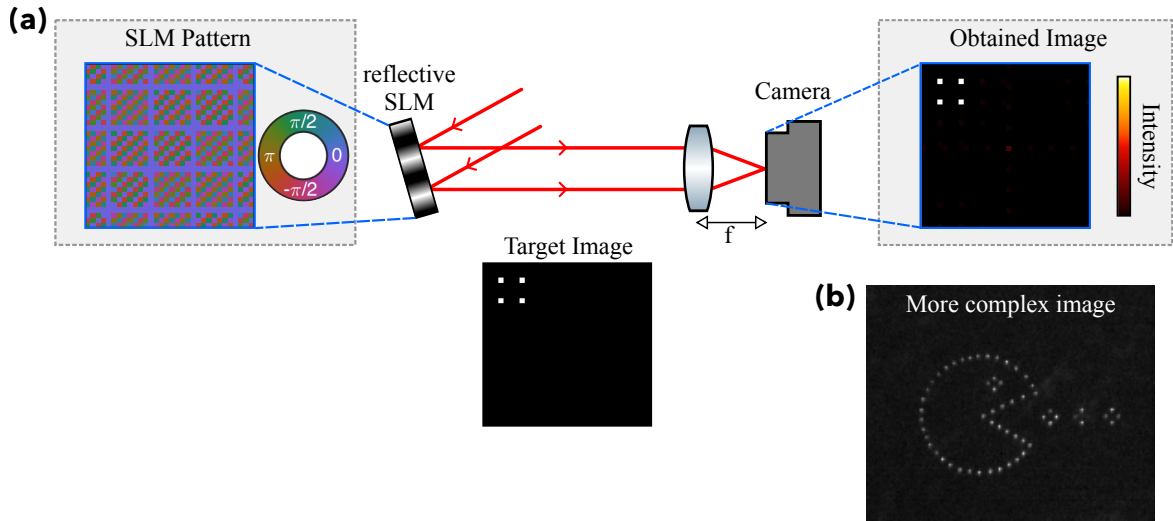


Figure C.1: Holographically generating a 2D array of trapped atoms. **(a)** A collimated laser beam picks up a phase pattern from the **SLM** and, once it passes through the lens, forms a 2D image in the focal plane of the lens. Each dot can correspond to one tweezer/trapped atom. Adapted from [108]. **(b)** We can generate more intricate trapping patterns, as this particular image of the fluorescence of single atoms shows. Reproduced from [49].

1064 nm fibre laser⁹ illuminates the **SLM**, and so the **SLM** imprints a phase pattern onto the beam. More precisely, the **SLM** encodes a hologram of the desired 2D trapping pattern onto the beam. Once it passes through the tweezer lens, then because the lens performs a Fourier transform on the impinging light [42], the lens transforms the phase pattern into the desired 2D array at the focal plane, such as in figure C.1b. If we want a different array or to move the atoms, we just update the **SLM** accordingly. Because the tweezer lenses are specified such that the focal points of 1064 nm and 780 nm light are at the same position, we can simultaneously send the tweezer beam into the cavity and measure fluorescence coming out from atoms in the cavity.¹⁰ Altogether, we can construct reconfigurable 2D arrays of atoms in tweezers, placing them with micrometre precision at the cavity mode.¹¹

⁹ Quantel EYLSA 1064

¹⁰ We can separate the tweezer light and fluorescence light by using a dichroic mirror.

¹¹ Our research group has also produced 2D arrays of tweezers non-holographically [106], but in that case producing tweezers is much less power efficient; we need more laser power so that we can produce larger arrays.

To summarise this section, we have elucidated the theory of how to construct a 2D array of single atoms using optical tweezers. For further information on optical tweezers/dipole force traps, see Chu et al. [228] and the reviews on tweezers for neutral atoms [229] or small objects in general [230].

Bibliography

- [1] R. P. Feynman, *Simulating Physics with Computers*, [International Journal of Theoretical Physics](#) **21** (1982) 467 (cit. on p. 1).
- [2] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin and X. Yuan, *Quantum computational chemistry*, [Reviews of Modern Physics](#) **92** (2020) 015003 (cit. on p. 1).
- [3] V. Vedral, *Introduction to Quantum Information Science*, Oxford University Press, 2006 (cit. on pp. 1, 3).
- [4] V. Mavroeidis, K. Vishi, D. Mateusz and A. Jøsang, *The Impact of Quantum Computing on Present Cryptography*, [International Journal of Advanced Computer Science and Applications](#) **9** (2018) (cit. on p. 1).
- [5] F. Arute et al., *Quantum supremacy using a programmable superconducting processor*, [Nature](#) **574** (2019) 505 (cit. on p. 1).
- [6] J. Preskill, *Quantum computing and the entanglement frontier*, (2012), arXiv: [1203.5813 \[quant-ph\]](#) (cit. on p. 2).
- [7] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis and H. Neven, *Characterizing quantum supremacy in near-term devices*, [Nature Physics](#) **14** (2018) 595 (cit. on p. 2).
- [8] Z. Chen et al., *Exponential suppression of bit or phase errors with cyclic error correction*, [Nature](#) **595** (2021) 383 (cit. on p. 2).
- [9] J. Preskill, *Quantum Computing in the NISQ era and beyond*, [Quantum](#) **2** (2018) 79 (cit. on p. 2).
- [10] D. P. Divincenzo, *The Physical Implementation of Quantum Computation*, [Fortschritte der Physik](#) **48** (2000) 771 (cit. on p. 2).
- [11] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe and J. L. O'Brien, *Quantum computers*, [Nature](#) **464** (2010) 45 (cit. on p. 2).

- [12] A. Fruchtman and I. Choi, *Technical Roadmap for Fault-Tolerant Quantum Computing*, tech. rep., University of Oxford, 2016,
URL: <https://nqit.ox.ac.uk/sites/www.nqit.ox.ac.uk/files/2016-11/NQIT%20Technical%20Roadmap.pdf> (cit. on p. 2).
- [13] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright and C. Monroe, *Demonstration of a small programmable quantum computer with atomic qubits*, *Nature* **536** (2016) 63 (cit. on pp. 2, 192).
- [14] H. J. Kimble, *The quantum internet*, *Nature* **453** (2008) 1023 (cit. on pp. 2, 44, 80, 205).
- [15] T. E. Northup and R. Blatt, *Quantum information transfer using photons*, *Nature Photonics* **8** (2014) 356 (cit. on pp. 2, 44).
- [16] C. Nölleke, A. Neuzner, A. Reiserer, C. Hahn, G. Rempe and S. Ritter, *Efficient Teleportation Between Remote Single-Atom Quantum Memories*, *Physical Review Letters* **110** (2013) 140403 (cit. on pp. 2, 79).
- [17] M. Uphoff, M. Brekenfeld, G. Rempe and S. Ritter, *An integrated quantum repeater at telecom wavelength with single atoms in optical fiber cavities*, *Applied Physics B* **122** (2016) 46 (cit. on pp. 2, 3).
- [18] C. Degen, F. Reinhard and P. Cappellaro, *Quantum sensing*, *Reviews of Modern Physics* **89** (2017) 035002 (cit. on p. 3).
- [19] X. Liu, X. Yao, R. Xue, H. Wang, H. Li, Z. Wang, L. You, X. Feng, F. Liu, K. Cui, Y. Huang and W. Zhang, *An entanglement-based quantum network based on symmetric dispersive optics quantum key distribution*, *APL Photonics* **5** (2020) 076104 (cit. on p. 3).
- [20] M. A. Nielsen, *Cluster-state quantum computation*, *Reports on Mathematical Physics* **57** (2006) 147 (cit. on p. 3).
- [21] A. Reiserer and G. Rempe, *Cavity-based quantum networks with single atoms and optical photons*, *Reviews of Modern Physics* **87** (2015) 1379 (cit. on pp. 3, 44, 76, 84, 88, 95, 192, 205).
- [22] T. Wilk, S. C. Webster, A. Kuhn and G. Rempe, *Single-Atom Single-Photon Quantum Interface*, *Science* **317** (2007) 488 (cit. on p. 3).
- [23] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann and G. Rempe, *An elementary quantum network of single atoms in optical cavities*, *Nature* **484** (2012) 195 (cit. on pp. 3, 44, 73, 79, 95).

Bibliography

- [24] S. Daiss, S. Langenfeld, S. Welte, E. Distanto, P. Thomas, L. Hartung, O. Morin and G. Rempe, *A quantum-logic gate between distant quantum-network modules*, *Science* **371** (2021) 614 (cit. on pp. 3, 74, 195).
- [25] J. I. Cirac, P. Zoller, H. J. Kimble and H. Mabuchi, *Quantum State Transfer and Entanglement Distribution among Distant Nodes in a Quantum Network*, *Physical Review Letters* **78** (1997) 3221 (cit. on pp. 3, 40, 49, 56, 75).
- [26] A. Kuhn, “Cavity Induced Interfacing of Atoms and Light”, *Engineering the Atom-Photon Interaction*, Springer International Publishing, 2015, chap. 1 (cit. on pp. 3, 67).
- [27] J. Volz, R. Gehr, G. Dubois, J. Estève and J. Reichel, *Measurement of the internal state of a single atom without energy exchange*, *Nature* **475** (2011) 210 (cit. on pp. 3, 195).
- [28] A. Reiserer, S. Ritter and G. Rempe, *Nondestructive Detection of an Optical Photon*, *Science* **342** (2013) 1349 (cit. on p. 3).
- [29] D. Niemietz, P. Farrera, S. Langenfeld and G. Rempe, *Nondestructive detection of photonic qubits*, *Nature* **591** (2021) 570 (cit. on pp. 3, 195, 199).
- [30] J. Hofmann, M. Krug, N. Ortegel, L. Gérard, M. Weber, W. Rosenfeld and H. Weinfurter, *Heralded Entanglement Between Widely Separated Atoms*, *Science* **337** (2012) 72 (cit. on p. 4).
- [31] T. D. Barrett, A. Rubenok, D. Stuart, O. Barter, A. Holleczek, J. Dille, P. B. R. Nisbet-Jones, K. Poulios, G. D. Marshall, J. L. O’Brien, A. Politi, J. C. F. Matthews and A. Kuhn, *Multimode interferometry for entangling atoms in quantum networks*, *Quantum Science and Technology* **4** (2019) 025008 (cit. on p. 4).
- [32] T. D. Barrett, O. Barter, D. Stuart, B. Yuen and A. Kuhn, *Polarization Oscillations in Birefringent Emitter-Cavity Systems*, *Physical Review Letters* **122** (2019) 083602 (cit. on pp. 4, 50, 73, 77, 99).
- [33] S. Wehner, D. Elkouss and R. Hanson, *Quantum internet: A vision for the road ahead*, *Science* **362** (2018) eaam9288 (cit. on pp. 4, 44).
- [34] C. J. Hood, H. J. Kimble and J. Ye, *Characterization of high-finesse mirrors: Loss, phase shifts, and mode structure in an optical cavity*, *Physical Review A* **64** (2001) 033804 (cit. on pp. 4, 88, 90, 112).
- [35] A. Saharyan, J.-R. Álvarez, T. H. Doherty, A. Kuhn and S. Guérin, *Light-matter interaction in open cavities with dielectric stacks*, *Applied Physics Letters* **118** (2021) 154002 (cit. on p. 4).

Bibliography

- [36] A. Reiserer, C. Nölleke, S. Ritter and G. Rempe, *Ground-State Cooling of a Single Atom at the Center of an Optical Cavity*, *Physical Review Letters* **110** (2013) 223003 (cit. on p. 5).
- [37] T. Nugent, “The Oval from the air”, licensed under CC BY-SA 2.0, 2015, URL: <https://www.geograph.org.uk/photo/4418545> (cit. on p. 7).
- [38] E. M. Purcell, *Spontaneous Emission Probabilities at Radio Frequencies*, *Physical Review* **69** (1946) 674 (cit. on pp. 7, 11).
- [39] S. Haroche and J.-M. Raimond, *Exploring the Quantum*, Oxford University Press, 2006 (cit. on pp. 7, 54, 176, 204).
- [40] M. Fox, *Quantum Optics: An Introduction*, Oxford University Press, 2005 (cit. on pp. 11, 30–32).
- [41] N. Ismail, C. C. Kores, D. Geskus and M. Pollnau, *Fabry-Pérot resonator: spectral line shapes, generic and related Airy distributions, linewidths, finesses, and performance at low or frequency-dependent reflectivity*, *Optics Express* **24** (2016) 16366 (cit. on p. 11).
- [42] D. Meschede, *Optics, Light, and Lasers*, 3rd ed., Wiley-VCH Verlag GmbH & Co. KGaA, 2017 (cit. on pp. 12–15, 140, 164, 198, 219, 221).
- [43] A. E. Siegman, *Lasers*, University Science Books, 1986 (cit. on pp. 12, 17, 102).
- [44] H. Kogelnik and T. Li, *Laser Beams and Resonators*, *Applied Optics* **5** (1966) 1550 (cit. on p. 14).
- [45] G. Brooker, *Modern Classical Optics (Oxford Master Series in Physics)*, Oxford University Press, 2003 (cit. on pp. 14, 16).
- [46] M. Fischer, B. Srivathsan, L. Alber, M. Weber, M. Sondermann and G. Leuchs, *Shifting the phase of a coherent beam with a $^{174}\text{Yb}^+$ ion: influence of the scattering cross section*, *Applied Physics B* **123** (2017) (cit. on p. 15).
- [47] W. T. Silfvast, *Laser Fundamentals*, 2nd ed., Cambridge University Press, 2008 (cit. on p. 16).
- [48] R. Loudon, *The Quantum Theory of Light*, 3rd ed., Oxford University Press, 2000 (cit. on pp. 18, 20, 21).
- [49] N. Holland, *Holographically Generated Optical Tweezers for the Control of Single Atoms*, PhD thesis: University of Oxford, 2019 (cit. on pp. 20, 141, 146, 159, 220, 221).
- [50] V. Vedral, *Modern Foundations of Quantum Optics*, World Scientific, 2005 (cit. on pp. 21, 31, 37).
- [51] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, Cambridge University Press, 2017 (cit. on pp. 23, 215).

- [52] M. Oberst, H. Münch, G. Grigoryan and T. Halfmann, *Stark-chirped rapid adiabatic passage among a three-state molecular system: Experimental and numerical investigations*, *Physical Review A* **78** (2008) 033409 (cit. on p. 29).
- [53] A. Neuzner, M. Körber, S. Dürr, G. Rempe and S. Ritter, *Breakdown of atomic hyperfine coupling in a deep optical-dipole trap*, *Physical Review A* **92** (2015) (cit. on pp. 29, 220).
- [54] H. J. Kimble, *Strong Interactions of Single Atoms and Photons in Cavity QED*, *Physica Scripta* **T76** (1998) 127 (cit. on pp. 33, 205).
- [55] T. Barrett, *Polarised Single Photons from a Cavity-Enhanced Atom-Light Interface in Photonic Quantum Networks*, PhD thesis: University of Oxford, 2018, URL: <https://ora.ox.ac.uk/objects/uuid:c4263120-ca20-4fbd-9995-1eee6b097363> (cit. on pp. 36, 77, 78, 84, 132, 151, 194, 203).
- [56] E. Brion, L. H. Pedersen and K. Mølmer, *Adiabatic elimination in a lambda system*, *Journal of Physics A: Mathematical and Theoretical* **40** (2007) 1033 (cit. on p. 36).
- [57] N. V. Vitanov, A. A. Rangelov, B. W. Shore and K. Bergmann, *Stimulated Raman adiabatic passage in physics, chemistry, and beyond*, *Reviews of Modern Physics* **89** (2017) (cit. on pp. 39, 41).
- [58] B. W. Shore, *Picturing stimulated Raman adiabatic passage: a STIRAP tutorial*, *Advances in Optics and Photonics* **9** (2017) 563 (cit. on pp. 39, 40, 54).
- [59] J. Dilley, P. Nisbet-Jones, B. W. Shore and A. Kuhn, *Single-photon absorption in coupled atom-cavity systems*, *Physical Review A* **85** (2012) 023834 (cit. on pp. 40, 45, 56, 69, 79).
- [60] H. P. Specht, C. Nölleke, A. Reiserer, M. Uphoff, E. Figueroa, S. Ritter and G. Rempe, *A single-atom quantum memory*, *Nature* **473** (2011) 190 (cit. on p. 40).
- [61] T. Macha, E. Uruñuela, W. Alt, M. Ammenwerth, D. Pandey, H. Pfeifer and D. Meschede, *Nonadiabatic storage of short light pulses in an atom-cavity system*, *Physical Review A* **101** (2020) (cit. on pp. 40, 47, 48, 68, 72).
- [62] A. Kuhn, M. Hennrich, T. Bundo and G. Rempe, *Controlled generation of single photons from a strongly coupled atom-cavity system*, *Applied Physics B* **69** (1999) 373 (cit. on pp. 41, 54, 55).
- [63] B. Casabone, K. Friebe, B. Brandstätter, K. Schüppert, R. Blatt and T. Northup, *Enhanced Quantum Interface with Collective Ion-Cavity Coupling*, *Physical Review Letters* **114** (2015) 023602 (cit. on p. 44).
- [64] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, *Science* **356** (2017) 928 (cit. on pp. 44, 193).

- [65] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham and R. Hanson, *Deterministic delivery of remote entanglement on a quantum network*, *Nature* **558** (2018) 268 (cit. on pp. 44, 45, 79, 193).
- [66] H. Takahashi, E. Kassa, C. Christoforou and M. Keller, *Strong Coupling of a Single Ion to an Optical Cavity*, *Physical Review Letters* **124** (2020) (cit. on pp. 44, 82, 192, 196).
- [67] L. Stephenson, D. Nadlinger, B. Nichol, S. An, P. Drmota, T. Ballance, K. Thirumalai, J. Goodwin, D. Lucas and C. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network*, *Physical Review Letters* **124** (2020) 110501 (cit. on pp. 45, 192).
- [68] E. Janitz, M. K. Bhaskar and L. Childress, *Cavity quantum electrodynamics with color centers in diamond*, *Optica* **7** (2020) 1232 (cit. on p. 45).
- [69] P. Michler, *Quantum Dots for Quantum Information Technologies*, Springer International Publishing, 2017 (cit. on pp. 45, 192).
- [70] D. Lago-Rivera, S. Grandi, J. V. Rakonjac, A. Seri and H. de Riedmatten, *Telecom-heralded entanglement between multimode solid-state quantum memories*, *Nature* **594** (2021) 37 (cit. on p. 45).
- [71] X. Liu, J. Hu, Z.-F. Li, X. Li, P.-Y. Li, P.-J. Liang, Z.-Q. Zhou, C.-F. Li and G.-C. Guo, *Heralded entanglement distribution between two absorptive quantum memories*, *Nature* **594** (2021) 41 (cit. on p. 45).
- [72] K. T. Kaczmarek, P. M. Ledingham, B. Brecht, S. E. Thomas, G. S. Thekkadath, O. Lazo-Arjona, J. H. D. Munns, E. Poem, A. Feizpour, D. J. Saunders, J. Nunn and I. A. Walmsley, *High-speed noise-free optical quantum memory*, *Physical Review A* **97** (2018) 042316 (cit. on p. 45).
- [73] J. A. Nelder and R. Mead, *A Simplex Method for Function Minimization*, *The Computer Journal* **7** (1965) 308 (cit. on pp. 45, 50, 51).
- [74] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., Springer New York, 2006 (cit. on pp. 45, 50).
- [75] G. S. Vasilev, D. Ljunggren and A. Kuhn, *Single photons made-to-measure*, *New Journal of Physics* **12** (2010) 063024 (cit. on pp. 45, 69, 79).
- [76] A. V. Gorshkov, A. André, M. Fleischhauer, A. S. Sørensen and M. D. Lukin, *Universal approach to optimal photon storage in atomic media*, *Physical Review Letters* **98** (2007) (cit. on pp. 46, 69, 79).

- [77] J. L. Sørensen, D. Møller, T. Iversen, J. B. Thomsen, F. Jensen, P. Staunum, D. Voigt and M. Drewsen, *Efficient coherent internal state transfer in trapped ions using stimulated Raman adiabatic passage*, *New Journal of Physics* **8** (2006) 261 (cit. on p. 46).
- [78] M. Mücke, J. Bochmann, C. Hahn, A. Neuzner, C. Nölleke, A. Reiserer, G. Rempe and S. Ritter, *Generation of single photons from an atom-cavity system*, *Physical Review A* **87** (2013) 063805 (cit. on p. 46).
- [79] N. Maring, P. Farrera, K. Kutluer, M. Mazzer, G. Heinze and H. de Riedmatten, *Photonic quantum state transfer between a cold atomic gas and a crystal*, *Nature* **551** (2017) 485 (cit. on p. 46).
- [80] X.-L. Pang, A.-L. Yang, J.-P. Dou, H. Li, C.-N. Zhang, E. Poem, D. J. Saunders, H. Tang, J. Nunn, I. A. Walmsley and X.-M. Jin, *A hybrid quantum memory-enabled network at room temperature*, *Science Advances* **6** (2020) eaax1425 (cit. on p. 46).
- [81] M. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, D. M. Reich and C. Koch, *Krotov: A Python implementation of Krotov's method for quantum optimal control*, *SciPost Physics* **7** (2019) (cit. on pp. 46, 57, 58, 60–62).
- [82] S. E. Sklarz and D. J. Tannor, *Loading a Bose-Einstein condensate onto an optical lattice: An application of optimal control theory to the nonlinear Schrödinger equation*, *Physical Review A* **66** (2002) (cit. on pp. 46, 57, 60).
- [83] D. Manzano, *A short introduction to the Lindblad master equation*, *AIP Advances* **10** (2020) 025106 (cit. on pp. 47, 62, 215).
- [84] D. F. Walls and G. J. Milburn, *Quantum Optics*, 2nd ed., Springer Berlin Heidelberg, 2008 (cit. on pp. 47, 48).
- [85] T. D. Barrett, *rb-cqed*, GitHub repository, 2019, URL: <https://github.com/tomdbar/rb-cqed> (cit. on p. 49).
- [86] L. Han and M. Neumann, *Effect of dimensionality on the Nelder–Mead simplex method*, *Optimization Methods and Software* **21** (2006) 1 (cit. on p. 51).
- [87] A. Blaessle, *constrNMPy*, GitHub repository, 2018, URL: <https://github.com/alexblaessle/constrNMPy> (cit. on p. 52).
- [88] D. Steck, *Alkali D line data*, 2013, URL: <http://steck.us/alkalidata/> (cit. on pp. 53, 134, 138, 146, 160, 161).
- [89] C. P. Koch, *Controlling open quantum systems: tools, achievements, and limitations*, *Journal of Physics: Condensed Matter* **28** (2016) 213001 (cit. on p. 57).

-
- [90] F. K. Wilhelm, S. Kirchhoff, S. Machnes, N. Wittler and D. Sugny, *An introduction into optimal control for quantum technologies*, (2020), arXiv: [2003.10132 \[quant-ph\]](https://arxiv.org/abs/2003.10132) (cit. on p. 57).
- [91] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen and S. J. Glaser, *Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms*, *Journal of Magnetic Resonance* **172** (2005) 296 (cit. on p. 57).
- [92] D. M. Reich, M. Ndong and C. P. Koch, *Monotonically convergent optimization in quantum control using Krotov's method*, *The Journal of Chemical Physics* **136** (2012) 104103 (cit. on pp. 57, 60).
- [93] L. Giannelli, T. Schmit, T. Calarco, C. P. Koch, S. Ritter and G. Morigi, *Optimal storage of a single photon by a single intra-cavity atom*, *New Journal of Physics* **20** (2018) 105009 (cit. on pp. 64, 72).
- [94] P. Kumar, S. A. Malinovskaya and V. S. Malinovsky, *Optimal control of population and coherence in three-level Λ systems*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **44** (2011) 154010 (cit. on p. 64).
- [95] M. H. Goerz and K. Jacobs, *Efficient optimization of state preparation in quantum networks using quantum trajectories*, *Quantum Science and Technology* **3** (2018) 045005 (cit. on p. 65).
- [96] T. Macha, *Storage of short light pulses in a fiber-based atom-cavity system*, PhD thesis: Universität Bonn, 2018, URL: <https://hdl.handle.net/20.500.11811/7899> (cit. on p. 66).
- [97] L. N. Trefethen, *Approximation Theory and Approximation Practice, Extended Edition*, Society for Industrial and Applied Mathematics, 2019 (cit. on p. 68).
- [98] O. Morin, M. Körber, S. Langenfeld and G. Rempe, *Deterministic Shaping and Reshaping of Single-Photon Temporal Wave Functions*, *Physical Review Letters* **123** (2019) 133602 (cit. on p. 69).
- [99] T. Caneva, T. Calarco and S. Montangero, *Chopped random-basis quantum optimization*, *Physical Review A* **84** (2011) 022326 (cit. on p. 73).
- [100] P. Doria, T. Calarco and S. Montangero, *Optimal Control Technique for Many-Body Quantum Dynamics*, *Physical Review Letters* **106** (2011) 190501 (cit. on p. 73).
- [101] T. D. Barrett, D. Stuart, O. Barter and A. Kuhn, *Nonlinear Zeeman effects in the cavity-enhanced emission of polarised photons*, *New Journal of Physics* **20** (2018) 073030 (cit. on p. 73).

- [102] T. D. Barrett, T. H. Doherty and A. Kuhn, *Pushing Purcell enhancement beyond its limits*, [New Journal of Physics](#) **22** (2020) 063013 (cit. on pp. 73, 99, 196).
- [103] B. Vogell, B. Vermersch, T. E. Northup, B. P. Lanyon and C. A. Muschik, *Deterministic quantum state transfer between remote qubits in cavities*, [Quantum Science and Technology](#) **2** (2017) 045003 (cit. on p. 74).
- [104] P. Loock, W. Alt, C. Becher, O. Benson, H. Boche, C. Deppe, J. Eschner, S. Höfling, D. Meschede, P. Michler, F. Schmidt and H. Weinfurter, *Extending Quantum Links: Modules for Fiber- and Memory-Based Quantum Repeaters*, [Advanced Quantum Technologies](#) **3** (2020) 1900141 (cit. on p. 74).
- [105] T. M. Barlow, R. Bennett and A. Beige, *A master equation for a two-sided optical cavity*, [Journal of Modern Optics](#) **62** (2015) S11 (cit. on pp. 76, 198, 200).
- [106] C. Muldoon, L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsforth and A. Kuhn, *Control and manipulation of cold atoms in optical tweezers*, [New Journal of Physics](#) **14** (2012) 073051 (cit. on pp. 82, 195, 221).
- [107] D. Stuart and A. Kuhn, *Single-atom trapping and transport in DMD-controlled optical tweezers*, [New Journal of Physics](#) **20** (2018) 023013 (cit. on pp. 82, 141, 220).
- [108] N. Holland, D. Stuart, O. Barter and A. Kuhn, *Benchmarking modern algorithms to holographically create optical tweezers for laser-cooled atoms*, [Journal of Modern Optics](#) **65** (2018) 2133 (cit. on pp. 82, 141, 195, 220, 221).
- [109] T. G. Ballance, H. M. Meyer, P. Kobel, K. Ott, J. Reichel and M. Köhl, *Cavity-induced backaction in Purcell-enhanced photon emission of a single ion in an ultraviolet fiber cavity*, [Physical Review A](#) **95** (2017) 033812 (cit. on pp. 82, 192, 196).
- [110] M. Brekenfeld, D. Niemietz, J. D. Christesen and G. Rempe, *A quantum network node with crossed optical fibre cavities*, [Nature Physics](#) **16** (2020) 647 (cit. on pp. 82, 195, 196, 199).
- [111] S. Bogdanović, S. B. van Dam, C. Bonato, L. C. Coenen, A.-M. J. Zwerver, B. Hensen, M. S. Z. Liddy, T. Fink, A. Reiserer, M. Lončar and R. Hanson, *Design and low-temperature characterization of a tunable microcavity for diamond-based quantum networks*, [Applied Physics Letters](#) **110** (2017) 171103 (cit. on p. 82).
- [112] S. B. van Dam, M. Ruf and R. Hanson, *Optimal design of diamond-air microcavities for quantum networks using an analytical approach*, [New Journal of Physics](#) **20** (2018) 115004 (cit. on pp. 82, 194).

Bibliography

- [113] E. Kassa, H. Takahashi, C. Christoforou and M. Keller, *Precise positioning of an ion in an integrated Paul trap-cavity system using radiofrequency signals*, *Journal of Modern Optics* **65** (2017) 520 (cit. on pp. 82, 192, 196).
- [114] E. Kyoseva, A. Beige and L. C. Kwek, *Coherent cavity networks with complete connectivity*, *New Journal of Physics* **14** (2012) 023023 (cit. on p. 82).
- [115] J. Benedikter, T. Hümmer, M. Mader, B. Schlederer, J. Reichel, T. W. Hänsch and D. Hunger, *Transverse-mode coupling and diffraction loss in tunable Fabry–Pérot microcavities*, *New Journal of Physics* **17** (2015) 053051 (cit. on pp. 83, 96).
- [116] N. Podoliak, H. Takahashi, M. Keller and P. Horak, *Harnessing the mode mixing in optical fiber-tip cavities*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **50**, 085503 (2017) 085503 (cit. on p. 83).
- [117] T. Steinmetz, Y. Colombe, D. Hunger, T. W. Hänsch, A. Balocchi, R. J. Warburton and J. Reichel, *Stable fiber-based Fabry–Pérot cavity*, *Applied Physics Letters* **89** (2006) 111110 (cit. on p. 84).
- [118] D. Hunger, T. Steinmetz, Y. Colombe, C. Deutsch, T. W. Hänsch and J. Reichel, *A fiber Fabry–Perot cavity with high finesse*, *New Journal of Physics* **12** (2010) 065038 (cit. on pp. 84, 86, 93, 97, 101).
- [119] G. Rempe, *Atoms In An Optical Cavity - Quantum Electrodynamics In Confined Space*, *Contemporary Physics* **34** (1993) 119 (cit. on p. 84).
- [120] R. Miller, T. E. Northup, K. M. Birnbaum, A. Boca, A. D. Boozer and H. J. Kimble, *Trapped atoms in cavity QED: coupling quantized light and matter*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **38** (2005) S551 (cit. on p. 84).
- [121] D. Kleckner, W. T. M. Irvine, S. S. R. Oemrawsingh and D. Bouwmeester, *Diffraction-limited high-finesse optical cavities*, *Physical Review A* **81** (2010) 043814 (cit. on p. 84).
- [122] A. A. P. Trichet, P. R. Dolan, D. M. Coles, G. M. Hughes and J. M. Smith, *Topographic control of open-access microcavities at the nanometer scale*, *Optics Express* **23** (2015) 17205 (cit. on pp. 84, 86, 196).
- [123] G. Rempe, R. Lalezari, R. J. Thompson and H. J. Kimble, *Measurement of ultralow losses in an optical interferometer*, *Optics Letters* **17** (1992) 363 (cit. on p. 84).

- [124] T. E. Northup, *Coherent control in cavity QED*, PhD thesis: California Institute of Technology, 2008, URL: <http://resolver.caltech.edu/CaltechETD:etd-05242008-114227> (cit. on p. 84).
- [125] J. Dong, *Laser trapping of atoms and cavity quantum electrodynamics in fibre-tip microcavities*, PhD thesis: University of Oxford, 2014, URL: <https://ora.ox.ac.uk/objects/uuid:1198a4d0-107b-4e07-87fc-238379d0840f> (cit. on pp. 84, 85, 97, 99, 101, 106, 128, 132).
- [126] C. Deutsch, *High Finesse Fibre Fabry-Perot Resonators - Production, Characterization, and Applications*, Diploma thesis: Ludwig-Maximilians-Universität München, 2008, URL: <http://hdl.handle.net/11858/00-001M-0000-000F-B4D4-D> (cit. on pp. 85, 97, 101).
- [127] C. J. R. Sheppard and S. Saghafi, *Beam modes beyond the paraxial approximation: A scalar treatment*, *Physical Review A* **57** (1998) 2971 (cit. on p. 86).
- [128] L. Flatten, A. Trichet and J. Smith, *Spectral engineering of coupled open-access microcavities*, *Laser & Photonics Reviews* **10** (2015) 257 (cit. on p. 86).
- [129] J. Gallego, S. Ghosh, S. K. Alavi, W. Alt, M. Martinez-Dorantes, D. Meschede and L. Ratschbacher, *High-finesse fiber Fabry-Perot cavities: stabilization and mode matching analysis*, *Applied Physics B: Lasers and Optics* **122** (2016) 47 (cit. on pp. 88–90, 92, 96, 114, 115, 128, 129, 197, 203).
- [130] W. B. Joyce and B. C. DeLoach, *Alignment of Gaussian beams*, *Applied Optics* **23** (1984) 4187 (cit. on pp. 88–90, 93).
- [131] J. C. Gallego, *Strong Coupling between Small Atomic Ensembles and an Open Fiber Cavity*, PhD thesis: Universität Bonn, 2017, URL: <https://hdl.handle.net/20.500.11811/7530> (cit. on pp. 89, 96, 116, 120).
- [132] K. Ott, S. Garcia, R. Kohlhaas, K. Schüppert, P. Rosenbusch, R. Long and J. Reichel, *Millimeter-long fiber Fabry-Perot cavities*, *Optics Express* **24** (2016) 9839 (cit. on pp. 96, 101, 102, 196).
- [133] P. Eaton and P. West, *Atomic Force Microscopy*, Oxford University Press, 2010 (cit. on p. 97).
- [134] M. Uphoff, M. Brekenfeld, G. Rempe and S. Ritter, *Frequency splitting of polarization eigenmodes in microscopic Fabry-Perot cavities*, *New Journal of Physics* **17** (2015) 013053 (cit. on pp. 98, 99, 196).

- [135] H. Takahashi, J. Morphew, F. Oručević, A. Noguchi, E. Kassa and M. Keller, *Novel laser machining of optical fibers for long cavities with low birefringence*, *Optics Express* **22** (2014) 31317 (cit. on pp. 102, 196).
- [136] L. Henke, N. Nagy and U. J. Krull, *An AFM determination of the effects on surface roughness caused by cleaning of fused silica and glass substrates in the process of optical biosensor preparation*, *Biosensors and Bioelectronics* **17** (2002) 547 (cit. on p. 108).
- [137] J. P. Hamilton, A. K. Rickertsen, D. P. Hamilton and D. A. Giesen, “Elimination of Mirror Recoating and Realignment: Optical Contamination Control Using First Contact Polymers on GTC, Keck, LIGO and Starshade Telescopes (Conference Presentation)”, *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation III*, ed. by R. Geyl and R. Navarro, SPIE, 2018 (cit. on p. 108).
- [138] L. Yao, T. A. Birks and J. C. Knight, *Low bend loss in tightly-bent fibers through adiabatic bend transitions*, *Optics Express* **17** (2009) 2962 (cit. on p. 112).
- [139] T. G. Ballance, *An ultraviolet fibre-cavity for strong ion-photon interaction*, PhD thesis: University of Cambridge, 2017, URL: <https://doi.org/10.17863/CAM.13621> (cit. on pp. 116, 120, 123, 155).
- [140] F. Ferri, *Strong coupling between a fiber-cavity mode and a commensurate atom lattice: a new platform for many-body entanglement with single-particle control*, PhD thesis: École normale supérieure, 2018 (cit. on pp. 116, 120, 129, 195).
- [141] C. Saavedra, D. Pandey, W. Alt, H. Pfeifer and D. Meschede, *Tunable fiber Fabry-Perot cavities with high passive stability*, *Optics Express* **29** (2021) 974 (cit. on pp. 116, 197).
- [142] B. Ellis, W. R. Ashcroft, S. J. Shaw, W. J. Cantwell, H. H. Kausch, G. P. Johari, F. R. Jones and X. M. Chen, *Chemistry and Technology of Epoxy Resins*, 2nd ed., Springer Netherlands, 1993 (cit. on p. 123).
- [143] M. Mader, J. Reichel, T. W. Hänsch and D. Hunger, *A scanning cavity microscope*, *Nature Communications* **6**, 7249 (2015) 7249 (cit. on p. 126).
- [144] J. F. S. Brachmann, H. Kaupp, T. W. Hänsch and D. Hunger, *Photothermal effects in ultra-precisely stabilized tunable microcavities*, *Optics Express* **24** (2016) 21205 (cit. on p. 128).
- [145] P. Nisbet-Jones, *Shaping single photons*, PhD thesis: University of Oxford, 2012, URL: <https://ora.ox.ac.uk/objects/uuid:c75d4896-c5a8-42b8-a166-ffcd4166fc09> (cit. on p. 132).

Bibliography

- [146] A. Kuhn, M. Hennrich and G. Rempe, *Deterministic Single-Photon Source for Distributed Quantum Networking*, *Physical Review Letters* **89** (2002) 067901 (cit. on p. 137).
- [147] M. Fox, *A Student's Guide to Atomic Physics*, Cambridge University Press, 2018 (cit. on pp. 143, 160, 179, 193, 218).
- [148] S. Hooker and C. Webb, *Laser Physics*, Oxford University Press, 2010, 586 (cit. on pp. 144, 162).
- [149] P. A. Tipler and G. P. Mosca, *Physics for Scientists and Engineers*, 6th ed., Macmillan Education, 2007 (cit. on p. 145).
- [150] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*, Springer New York, 1999 (cit. on pp. 146, 147, 182, 188).
- [151] J. Dalibard and C. Cohen-Tannoudji, *Laser cooling below the Doppler limit by polarization gradients: simple theoretical models*, *Journal of the Optical Society of America B* **6** (1989) 2023 (cit. on p. 147).
- [152] G. Grynberg, A. Aspect, C. Fabre and C. Cohen-Tannoudji, *Introduction to Quantum Optics*, Cambridge University Press, 2009 (cit. on p. 148).
- [153] D. Stuart, *Manipulating Single Atoms with Optical Tweezers*, PhD thesis: University of Oxford, 2014,
URL: <https://ora.ox.ac.uk/objects/uuid:ab99e851-3c66-4688-8725-b7d1588c5db0> (cit. on pp. 149, 151, 182, 183, 220).
- [154] H. Youk, *Numerical study of quadrupole magnetic traps for neutral atoms: anti-Helmholtz coils and a U-chip*, *Canadian Undergraduate Physics Journal* **3** (2005) 13 (cit. on p. 151).
- [155] J. Dilley, *A single-photon source for quantum networking*, PhD thesis: University of Oxford, 2012,
URL: <https://ora.ox.ac.uk/objects/uuid:380a4aaf-e809-4fff-84c7-5b6a0856a6cf> (cit. on pp. 152, 184).
- [156] E. R. Abraham and E. A. Cornell, *Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems*, *Applied Optics* **37** (1998) 1762 (cit. on pp. 156, 157).
- [157] B. Buchholz and V. Ebert, *Compact, compression-free, displaceable, and resealable vacuum feedthrough with built-in strain relief for sensitive components such as optical fibers*, *Review of Scientific Instruments* **85** (2014) 055109 (cit. on p. 156).
- [158] A. R. Peerzada, C. M. Jobson, E. Kassa, J. Morphew, X. Fernandez-Gonzalvo and M. Keller, *Versatile optical fiber feedthroughs for ultra-high vacuum applications*, *Vacuum* **180** (2020) 109542 (cit. on p. 156).

- [159] D. L. Miller and N. T. Moshegov, *All-metal ultrahigh vacuum optical fiber feedthrough*, *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films* **19** (2001) 386 (cit. on p. 157).
- [160] K. M. Birnbaum, *Cavity QED with multilevel atoms*, PhD thesis, 2005, URL: <https://resolver.caltech.edu/CaltechETD:etd-05272005-103306> (cit. on p. 158).
- [161] E. D. Black, *An introduction to Pound–Drever–Hall laser frequency stabilization*, *American Journal of Physics* **69** (2001) 79 (cit. on pp. 163, 164, 168).
- [162] D. J. McCarron, S. A. King and S. L. Cornish, *Modulation transfer spectroscopy in atomic rubidium*, *Measurement Science and Technology* **19** (2008) 105601 (cit. on p. 164).
- [163] R. K. Raj, D. Bloch, J. J. Snyder, G. Camy and M. Ducloy, *High-Frequency Optically Heterodyned Saturation Spectroscopy Via Resonant Degenerate Four-Wave Mixing*, *Physical Review Letters* **44** (1980) 1251 (cit. on p. 164).
- [164] D. A. Smith and I. G. Hughes, *The role of hyperfine pumping in multilevel systems exhibiting saturated absorption*, *American Journal of Physics* **72** (2004) 631 (cit. on p. 165).
- [165] G. Puentes, *Laser frequency offset locking scheme for high-field imaging of cold atoms*, *Applied Physics B* **107** (2012) 11 (cit. on p. 166).
- [166] R. D. Blevins, *Formulas for Dynamics, Acoustics and Vibration*, John Wiley & Sons, 2016 (cit. on p. 172).
- [167] F. Bernardot, P. Nussenzevig, M. Brune, J. M. Raimond and S. Haroche, *Vacuum Rabi Splitting Observed on a Microscopic Atomic Sample in a Microwave Cavity*, *Europhysics Letters (EPL)* **17** (1992) 33 (cit. on p. 176).
- [168] W. Wohlleben, F. Chevy, K. Madison and J. Dalibard, *An atom faucet*, *The European Physical Journal D* **15** (2001) 237 (cit. on p. 177).
- [169] L. Cacciapuoti, A. Castrillo, M. de Angelis and G. Tino, *A continuous cold atomic beam from a magneto-optical trap*, *The European Physical Journal D* **15** (2001) 245 (cit. on p. 177).
- [170] J. M. Raimond, P. Goy, M. Gross, C. Fabre and S. Haroche, *Statistics of Millimeter-Wave Photons Emitted by a Rydberg-Atom Maser: An Experimental Study of Fluctuations in Single-Mode Superradiance*, *Physical Review Letters* **49** (1982) 1924 (cit. on p. 181).
- [171] Y.-C. Liu, B.-B. Li and Y.-F. Xiao, *Electromagnetically induced transparency in optical microcavities*, *Nanophotonics* **6** (2017) 789 (cit. on p. 181).

- [172] I. G. Hughes and T. P. A. Hase, *Measurements and their Uncertainties: A practical guide to modern error analysis*, Oxford University Press, 2010 (cit. on p. 186).
- [173] G. L. Gattobigio, T. Pohl, G. Labeyrie and R. Kaiser, *Scaling laws for large magneto-optical traps*, *Physica Scripta* **81** (2010) 025301 (cit. on p. 188).
- [174] J. Gallego, W. Alt, T. Macha, M. Martinez-Dorantes, D. Pandey and D. Meschede, *Strong Purcell Effect on a Neutral Atom Trapped in an Open Fiber Cavity*, *Physical Review Letters* **121** (2018) (cit. on pp. 189, 194, 196, 203).
- [175] S. Begley, M. Vogt, G. K. Gulati, H. Takahashi and M. Keller, *Optimized Multi-Ion Cavity Coupling*, *Physical Review Letters* **116** (2016) 223001 (cit. on p. 192).
- [176] C. Monroe and J. Kim, *Scaling the Ion Trap Quantum Processor*, *Science* **339** (2013) 1164 (cit. on p. 192).
- [177] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan and J. Kim, *Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects*, *Physical Review A* **89** (2014) 022317 (cit. on pp. 192, 198).
- [178] S. Welte, B. Hacker, S. Daiss, S. Ritter and G. Rempe, *Photon-Mediated Quantum Gate between Two Neutral Atoms in an Optical Cavity*, *Physical Review X* **8** (2018) 011018 (cit. on pp. 192, 195).
- [179] M. Keller, B. Lange, K. Hayasaka, W. Lange and H. Walther, *Stable long-term coupling of a single ion to a cavity mode*, *Journal of Modern Optics* **54** (2007) 1607 (cit. on p. 192).
- [180] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J.-N. Zhang, L.-M. Duan, D. Yum and K. Kim, *Single-qubit quantum memory exceeding ten-minute coherence time*, *Nature Photonics* **11** (2017) 646 (cit. on p. 192).
- [181] J. P. Home, M. J. McDonnell, D. J. Szwer, B. C. Keitch, D. M. Lucas, D. N. Stacey and A. M. Steane, *Memory coherence of a sympathetically cooled trapped-ion qubit*, *Physical Review A* **79** (2009) 050305 (cit. on p. 192).
- [182] H. Takahashi, A. Wilson, A. Riley-Watson, F. Oručević, N. Seymour-Smith, M. Keller and W. Lange, *An integrated fiber trap for single-ion photonics*, *New Journal of Physics* **15** (2013) 053011 (cit. on p. 192).
- [183] E. Kassa, *Single ion coupled to a high-finesse optical fibre cavity for CQED in the strong coupling regime*, PhD thesis: University of Sussex, 2017, URL: <http://sro.sussex.ac.uk/69317/> (cit. on p. 192).

Bibliography

- [184] S. Fischbach, M. v. Helversen, M. Schmidt, A. Kaganskiy, R. Schmidt, A. Schliwa, T. Heindel, S. Rodt and S. Reitzenstein, *A deterministically fabricated spectrally-tunable quantum dot based single-photon source*, (2018), arXiv: [1805.10623v1](https://arxiv.org/abs/1805.10623v1) [[cond-mat.mes-hall](#)] (cit. on p. 193).
- [185] P.-I. Schneider, N. Srocka, S. Rodt, L. Zschiedrich, S. Reitzenstein and S. Burger, *Numerical optimization of the extraction efficiency of a quantum-dot based single-photon emitter into a single-mode fiber*, *Optics Express* **26** (2018) 8479 (cit. on p. 193).
- [186] F. Liu, A. J. Brash, J. O’Hara, L. M. P. P. Martins, C. L. Phillips, R. J. Coles, B. Royall, E. Clarke, C. Bentham, N. Prtljaga, I. E. Itskevich, L. R. Wilson, M. S. Skolnick and A. M. Fox, *High Purcell factor generation of indistinguishable on-chip single photons*, *Nature Nanotechnology* **13** (2018) 835 (cit. on p. 193).
- [187] K. Hennessy, A. Badolato, M. Winger, D. Gerace, M. Atatüre, S. Gulde, S. Fält, E. L. Hu and A. Imamoglu, *Quantum nature of a strongly coupled single quantum dot–cavity system*, *Nature* **445** (2007) 896 (cit. on p. 193).
- [188] R. Stockill, C. L. Gall, C. Matthiesen, L. Huthmacher, E. Clarke, M. Hugues and M. Atatüre, *Quantum dot spin coherence governed by a strained nuclear environment*, *Nature Communications* **7** (2016) (cit. on p. 193).
- [189] A. Dréau, A. Tchekorava, A. E. Mahdaoui, C. Bonato and R. Hanson, *Quantum Frequency Conversion of Single Photons from a Nitrogen-Vacancy Center in Diamond to Telecommunication Wavelengths*, *Physical Review Applied* **9** (2018) 064031 (cit. on p. 193).
- [190] D. D. Awschalom, R. Hanson, J. Wrachtrup and B. B. Zhou, *Quantum technologies with optically interfaced solid-state spins*, *Nature Photonics* **12** (2018) 516 (cit. on p. 193).
- [191] D. Riedel, I. Söllner, B. J. Shields, S. Starosielec, P. Appel, E. Neu, P. Maletinsky and R. J. Warburton, *Deterministic Enhancement of Coherent Photon Generation from a Nitrogen-Vacancy Center in Ultrapure Diamond*, *Physical Review X* **7** (2017) 031040 (cit. on p. 194).
- [192] S. Johnson, P. R. Dolan, T. Grange, A. A. P. Trichet, G. Hornecker, Y. C. Chen, L. Weng, G. M. Hughes, A. A. R. Watt, A. Auffèves and J. M. Smith, *Tunable cavity coupling of the zero phonon line of a nitrogen-vacancy defect in diamond*, *New Journal of Physics* **17** (2015) 122003 (cit. on p. 194).
- [193] S. Johnson, P. R. Dolan and J. M. Smith, *Diamond photonics for distributed quantum networks*, *Progress in Quantum Electronics* **55** (2017) 129 (cit. on p. 194).

- [194] Y. Chu et al., *Coherent Optical Transitions in Implanted Nitrogen Vacancy Centers*, *Nano Letters* **14** (2014) 1982 (cit. on p. 194).
- [195] J. Hütner, T. Hoinkes, M. Becker, M. Rothhardt, A. Rauschenbeutel and S. M. Skoff, *Nanofiber-based high-Q microresonator for cryogenic applications*, *Optics Express* **28** (2020) 3249 (cit. on p. 194).
- [196] E. Shahmoon, D. S. Wild, M. D. Lukin and S. F. Yelin, *Theory of cavity QED with 2D atomic arrays*, (2020), arXiv: 2006.01972 [quant-ph] (cit. on p. 194).
- [197] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond and C. Jurczak, *Quantum computing with neutral atoms*, *Quantum* **4** (2020) 327 (cit. on p. 194).
- [198] S. Garcia, F. Ferri, J. Reichel and R. Long, *Overlapping two standing waves in a microcavity for a multi-atom photon interface*, *Optics Express* **28** (2020) 15515 (cit. on p. 195).
- [199] F. Ferri, S. Garcia, M. Baghdad, J. Reichel and R. Long, *Mapping optical standing-waves of an open-access Fabry–Perot cavity with a tapered fiber*, *Review of Scientific Instruments* **91** (2020) 033104 (cit. on p. 195).
- [200] S. Langenfeld, O. Morin, M. Körber and G. Rempe, *A network-ready random-access qubits memory*, *npj Quantum Information* **6** (2020) (cit. on p. 195).
- [201] P. Kobel, M. Breyer and M. Köhl, *Deterministic spin-photon entanglement from a trapped ion in a fiber Fabry–Perot cavity*, *npj Quantum Information* **7** (2021) (cit. on p. 195).
- [202] S. Garcia, F. Ferri, K. Ott, J. Reichel and R. Long, *Dual-wavelength fiber Fabry-Perot cavities with engineered birefringence*, *Optics Express* **26** (2018) 22249 (cit. on p. 196).
- [203] D. Hunger, C. Deutsch, R. J. Barbour, R. J. Warburton and J. Reichel, *Laser micro-fabrication of concave, low-roughness features in silica*, *AIP Advances* **2** (2012) 012119 (cit. on p. 196).
- [204] T. Ruelle, M. Poggio and F. Braakman, *Optimized single-shot laser ablation of concave mirror templates on optical fibers*, *Applied Optics* **58** (2019) 3784 (cit. on p. 196).
- [205] R. A. Maruf and M. Bajcsy, *On-chip splicer for coupling light between photonic crystal and solid-core fibers*, *Applied Optics* **56** (2017) 4680 (cit. on p. 196).
- [206] G. K. Gulati, H. Takahashi, N. Podoliak, P. Horak and M. Keller, *Fiber cavities with integrated mode matching optics*, *Scientific Reports* **7** (2017) 5556 (cit. on p. 196).

- [207] J. Schmitz, H. M. Meyer and M. Köhl,
Ultraviolet Fabry-Perot cavity with stable finesse under ultrahigh vacuum conditions,
[Review of Scientific Instruments](#) **90** (2019) 063102 (cit. on p. 197).
- [208] H. Takahashi, E. Kassa, C. Christoforou and M. Keller,
Cavity-induced anticorrelated photon-emission rates of a single ion,
[Physical Review A](#) **96** (2 2017) 023824 (cit. on p. 197).
- [209] A. V. Rynbach, G. Schwartz, R. F. Spivey, J. Joseph, G. Vrijsen and J. Kim,
Design and characterization of an integrated surface ion trap and micromirror optical cavity, [Applied Optics](#) **56** (2017) 6511 (cit. on p. 197).
- [210] S. Ragg, C. Decaroli, T. Lutz and J. P. Home,
Segmented ion-trap fabrication using high precision stacked wafers,
[Review of Scientific Instruments](#) **90** (2019) 103203 (cit. on p. 197).
- [211] J. Kim et al.,
1100 x 1100 port MEMS-based optical crossconnect with 4-dB maximum loss,
[IEEE Photonics Technology Letters](#) **15** (2003) 1537 (cit. on p. 198).
- [212] J. Metz, M. Trupke and A. Beige,
Robust Entanglement through Macroscopic Quantum Jumps,
[Physical Review Letters](#) **97** (2006) 040503 (cit. on p. 200).
- [213] C. Marr, A. Beige and G. Rempe,
Entangled-state preparation via dissipation-assisted adiabatic passages,
[Physical Review A](#) **68** (2003) (cit. on p. 200).
- [214] J.-R. Álvarez, M. IJspeert, O. Barter, B. Yuen, T. D. Barrett, D. Stuart, J. Dilley, A. Holleczek and A. Kuhn, *How to administer an antidote to Schrödinger's cat*, (2021), arXiv: [2106.09705 \[quant-ph\]](#) (cit. on p. 200).
- [215] T. W. Hänsch, *Nobel Lecture: Passion for precision*,
[Reviews of Modern Physics](#) **78** (2006) 1297 (cit. on p. 204).
- [216] B. Abbott et al.,
Observation of Gravitational Waves from a Binary Black Hole Merger,
[Physical Review Letters](#) **116** (2016) 061102 (cit. on p. 204).
- [217] H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne and S. P. Vyatchanin,
Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics,
[Physical Review D](#) **65** (2001) 022002 (cit. on p. 205).
- [218] A. Ashkin and J. Dziedzic,
Optical trapping and manipulation of viruses and bacteria, [Science](#) **235** (1987) 1517 (cit. on p. 217).
- [219] A. Ashkin and J. M. Dziedzic, *Optical Levitation by Radiation Pressure*,
[Applied Physics Letters](#) **19** (1971) 283 (cit. on p. 217).

Bibliography

- [220] A. Ashkin, *Acceleration and Trapping of Particles by Radiation Pressure*, *Physical Review Letters* **24** (1970) 156 (cit. on p. 217).
- [221] A. Ashkin, J. M. Dziedzic, J. E. Bjorkholm and S. Chu, *Observation of a single-beam gradient force optical trap for dielectric particles*, *Optics Letters* **11** (1986) 288 (cit. on p. 217).
- [222] P. Kumar and A. K. Sarma, *Optical force on two-level atoms by few-cycle-pulse Gaussian laser fields beyond the rotating-wave approximation*, *Physical Review A* **84** (2011) (cit. on p. 218).
- [223] N. Schlosser, G. Reymond and P. Grangier, *Collisional Blockade in Microscopic Optical Dipole Traps*, *Physical Review Letters* **89** (2002) (cit. on p. 220).
- [224] A. Fuhrmanek, R. Bourgain, Y. R. P. Sortais and A. Browaeys, *Light-assisted collisions between a few cold atoms in a microscopic dipole trap*, *Physical Review A* **85** (2012) (cit. on p. 220).
- [225] Y. H. Fung and M. F. Andersen, *Efficient collisional blockade loading of a single atom into a tight microtrap*, *New Journal of Physics* **17** (2015) 073011 (cit. on p. 220).
- [226] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner and M. D. Lukin, *Atom-by-atom assembly of defect-free one-dimensional cold atom arrays*, *Science* **354** (2016) 1024 (cit. on p. 220).
- [227] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye and A. Browaeys, *Synthetic three-dimensional atomic structures assembled atom by atom*, *Nature* **561** (2018) 79 (cit. on p. 220).
- [228] S. Chu, J. E. Bjorkholm, A. Ashkin and A. Cable, *Experimental Observation of Optically Trapped Atoms*, *Physical Review Letters* **57** (1986) 314 (cit. on p. 222).
- [229] R. Grimm, M. Weidemüller and Y. B. Ovchinnikov, “Optical Dipole Traps for Neutral Atoms”, *Advances In Atomic, Molecular, and Optical Physics*, Elsevier, 2000 95 (cit. on p. 222).
- [230] A. Ashkin, *History of optical trapping and manipulation of small-neutral particle, atoms, and molecules*, *IEEE Journal of Selected Topics in Quantum Electronics* **6** (2000) 841 (cit. on p. 222).

Acknowledgements

Ever seen [the recovery Joan Mir made](#)? What a ride this has been.

Firstly, thanks to my supervisor, Axel Kuhn. Blessed with amazing intuition, he is unparalleled in breaking down complex problems in a simple way. In fact, one of the many tricks up his sleeve is to solve a problem by exclusively using the following three phrases: ‘For sure...’, ‘The catch is...’, and ‘Be cautious!’. These form the basis states $\{| \text{affirm} \rangle, | \text{doubt} \rangle, | \text{negate} \rangle\}$. Since he can combine them in arbitrary superpositions and entanglements and keep these all in memory, it turns out that Axel is a Turing-complete quantum computer. I myself am blessed to have encountered his magic, as it has made me a better person overall.

Thanks to my Atom-Photon Connection colleagues. Every single member has enlightened me in how to become a better researcher. Although the egalitarian in me wants to thank every individual exactly equally, I feel that I should mention those who were at my PhD’s ‘key moments’, to remind us what getting one is all about. Thanks to Thomas ‘Interplay’ Doherty for showing us why Scotland is so great; to Naomi ‘2k split’ Holland for being remarkably (and nonchalantly) inspiring; and to Juan ‘What you will see’ Alvarez and Mark ‘Quaker Oats’ IJspeert whose rigour and passion have timely reminded me what matters most in physics. I owe many thanks to the three postdocs with whom I worked with at certain times: Klara Theophilo, Ezra Kassa and Ben Yuen. I sincerely hope you three can eventually enjoy the fruits of my labour, although at the time of writing I’m not sure how. This thesis isn’t edible, for a start.

My group colleague Tom ‘Ghost Rider’ Barrett deserves a special mention. It is truly remarkable that even though we never worked together directly, he was essential to the success of this project on many levels. Coming from a dark place (a laptop with low screen brightness), I am grateful for the time spent discussing physics and programming (and the banter), and especially for ‘rb-cqed’, which gave me the jump-start I needed at the perfect time (and to which I built on extensively to create chapter 3). What a selfless person.

I want to thank the wider physics community too. I cannot possibly thank everyone who had an impact on me, but be assured that I have respect for each and every one of you. Again, special mentions and thanks to the ‘key moments’ people. To Richard Berry in biophysics: your whole group was so welcoming, and I enjoyed the project a lot. To Gerhard ‘These are the modes’¹ Rempe for his inspirational enthusiasm, amazing research, and for the most

¹ During a tour of our lab, looking at live video from a camera focused at a cavity mirror that was scanning through the higher-order transverse modes: ‘Look! These are the modes. You see them. The modes! They

Bibliography

memorable conference slide I have ever seen.² To all the academics in Leeds University, who gave me the education, freedom to explore physics, and confidence I needed to get to this stage in the first place.

Thanks to all of my friends. To all in St Hugh's and all the other colleges, to the amazing people dotted across the country, to the businesspeople, the sportspeople, the engineers, the folkies, the game developers...you see, they are from all walks of life, and I'm blessed to have met you all. I know I've been away, but I will try to get back in the swing of things. There will be more to the story, I'm sure. And, to all on the overlapping part of the physics/friend Venn diagram. In particular, thanks to Magda 'Sum' Szczykulska, my esteemed colleague and friend from whom I have learnt the most of anyone, in physics and in life. Thanks to her, I'm becoming who I want to be.

Although I never met him, thanks to the late Mark 'k-punk' Fisher for giving me the initial confidence to write with the passion I didn't know I had. Ironically, I don't yet have the confidence to say more on that, but rest assured I'll do right by you and your spirit.

Huge thanks to my family. To my older brother, Omar 'The Tank Engine', for the rock-solid support and advice on all matters, big and small. To my younger brother, Mahir, the most curious of us—you'll get your PhD soon! Looking forward to it. To my extended family in San Diego and elsewhere:³ I really appreciated how you went out of your way to make me feel welcome during my visits! Immense thanks to my parents, who supported me throughout. But most remarkably, when I ran out of funding, they took care of me and helped put a roof over my head so that I could focus fully on my work, rather than split myself between multiple jobs. I will be eternally grateful for that, and hope to repay you with my future success. Education really is important.

Finally, something small. I've done a lot to get to this point. I will allow myself this small victory. But, I have us all in mind. So, thank you all, and I will take some time to rest. Then, we will set off on a new journey together, underneath the blue skies ahead.

are there. These are the modes! Look! The modes, they are there!

² On the subject of quantum networks, showing an old depiction of Titanic sinking, complete with panicking women and children jumping into the sea: 'The problem is photon loss.'

³ Hi Grandmas! Look, I did it, I'm a doctor!

'Light is therefore colour.'

—J. M. W. Turner

Take a look at a £20 note.
Found the quote?

Colour is light.
Colour is quantum.
Nature is quantum.