Optical Cavities for Quantum Information Processing with Trapped Neutral Atoms

A thesis submitted for the degree of Doctor of Philosophy

Thomas H. Doherty

Exeter College

University of Oxford

Hilary 2021

In memory of Stephen Thomas Doherty (1955 - 2019)
and Thomas Joseph Doherty (1929 - 2020)

Optical Cavities for Quantum Information Processing with Trapped Neutral Atoms

Thomas Hamish Doherty, Exeter College
Department of Physics

Thesis submitted for the degree of Doctor of Philosophy
University of Oxford, Hilary 2021

Abstract

A leading strategy for the creation of a scalable quantum computer is through a networked architecture. This would comprise stationary nodes for quantum information processing and photonic channels for their communication and entanglement distribution. In this scheme a coherent interface between light and matter is a critical component, facilitating the deterministic and reversible transfer of quantum information between stationary and travelling qubits. This requirement may be fulfilled by a high finesse optical cavity, whose radiation field is strongly coupled to a trapped quantised emitter. Neutral atoms are an ideal species for the development of such a device, given their inherent compatibility with dielectric mirror surfaces. However, an outstanding challenge in building an elementary network of neutral atom-cavity systems is in establishing a single atom dipole trap within each cavity mode. This thesis presents the development of an open-access cavity suitable for single atom localisation, whilst maintaining strong atom-cavity coupling, $\{g_0, \kappa, \gamma\} = \{11, 2.1, 3\} \times 2\pi$MHz, and the efficient extraction of generated photons, $\eta = 0.5$. We consider cavity design methodology, deriving a connection between the physical properties of the optical resonator and its intended role in the production of single photons. An experimental platform is developed for the creation of cavity mirrors by laser ablation, with a high degree of geometrical control. This is used to produce a set of pyramidal micro-mirrors, demonstrated to represent a favourable alternative to the widespread adoption of fibre-tip cavities. These pyramidal mirrors are shown to support multi-feature resonators, allowing several atom-cavity interfaces to be formed using a pair of substrates. By developing a novel procedure for mirror alignment, a hybrid mirror cavity is fabricated and integrated into an experimental architecture tailored for atomic trapping. Suitability of the cavity for future experimentation is demonstrated by frequency stabilisation to an atomic resonance of ^{87}Rb.

Statement of Originality

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

The work of Chapter 7 involves the integration of an optical cavity with a single atom dipole trap. Atomic cooling and trapping has been a parallel line of research within our group. Efforts to ensure compatibility of the cavity with our developed platform of atomic trapping was performed in collaboration with group members, whose contributions are highlighted therein.

Thomas H. Doherty
Hilary 2021

Acknowledgements

First and foremost I would like to express my thanks to Axel Kuhn, for his introduction to the field of optical cavity fabrication and persistent efforts to contain my excitement and expenditure thereafter. I also thank our colleagues in the Networked Quantum Information Technologies Hub, Matthias Keller and Jason Smith, whose research groups graciously hosted me and openly shared their expertise.

I must give thanks to the Atom Photon Connection Group here in Oxford, both past and present, whose work laid the foundation for this thesis. To Ezra Kassa, without whom the pyramidal mirrors would not have been realised. To Tom Barrett, who shared my dream of an atom-cavity interface with perfect quantum efficiency. To Mark IJspeert, for his work in the unification of cavities with single atom dipole trapping, as well as providing a steady pair of hands in the laboratory when I was found lacking. To all the group members I was fortunate to share an office and research with; Juan, Marwan, Chloe, Ben, Oliver and Dustin.

Finally, to my family, Annie, Steve, Maureen and Tom, for their continual advice and support.

Contents

Chapter 1

Introduction

1.1 Networked Quantum Computation

Quantum mechanics emerged in the early 20th century as a framework to elucidate the nature of light and matter. Often described as a foundational work, Max Plank described the frequency spectrum of black body radiation via the quantisation of energy [1]. Then, by the division of this electromagnetic radiation into discrete wave-packets, or quanta, Albert Einstein made an account of the photoelectric effect [2]. The subsequent naming of these quantised electromagnetic field excitations was *photons* [3]. The intrinsic connection between atoms and photons was first highlighted in the explanation of spectral lines given by Niels Bohr, by applying quantisation to the hydrogen atom [4]. However, despite continuing significance, these theories are ultimately phenomenological in formulation and are today considered semi-classical.

Modern quantum mechanics considers atomic and photonic quantum systems to be inherently probabilistic; principally defined by predicting outcomes of their measurement. However, such probabilities are distinct from classical counterparts, as they do not represent uncertainty in the process of measurement, but the intrinsic state of the system in question. These probabilities obey non-intuitive restrictions, such as Heisenberg's uncertainty principle, limiting the maximum precision of their determination in connection to other measurable properties [5]. The process of measurement itself is inherently non-deterministic, applying discontinuous change to the state of the system seemingly instantaneously [6]. As one may expect, this led to extensive philosophical

discussion on the greater impact of quantum mechanics.

A core consequence of quantum mechanics is that a system may exist in a combination of states simultaneously, termed a *superposition*. These states may be highly correlated or *entangled* with a separate system, potentially over significant distances. Action on one system will simultaneously affect the other, demonstrating its nonlocality. In addition to the scientific interest this creates, it is being increasingly viewed as a potential technological resource. One example is quantum key distribution, providing a protocol for secure communication [7, 8]. Then, the field of quantum simulation overcomes the complexity of simulating quantum physics on a classical computer, by using an analogous quantum system with greater inherent control [9, 10].

A focal point of these emerging technologies is quantum information processing, where superposition and entanglement are used to facilitate computation [11]. These schemes are based on encoding information within quantum mechanical states, which in analogy to classical binary computation, are named qubits. Physically, a qubit is a coherently controllable two-state quantum system, existing in a superposition. Within a quantum computer, logical operations are performed on entangled qubits, establishing a platform quite unlike a classical computer. Certain tasks may now be performed with significantly less computational complexity, with the archetypal examples being algorithms for integer factorisation [12] and database searching [13].

However, the widespread anticipation for quantum computing must be measured against its outstanding technical challenges. To perform scalable and universal quantum computation, DiVincenzo outlined seven criteria that must be met [14]. This requires a system to have well defined qubits, whose state can be coherently controlled, initialised and modified by logical operations. The system must then be protected from decoherence until the computation is completed and its final state is read. A partial fulfilment of these requirements has been demonstrated using a large variety of different promising

platforms, including super-conducting circuits [15], quantum dots [16], diamond colour centres [17], photons [18], neutral atoms [19] and trapped ions [20].

A network of quantum information processing nodes is one pathway to achieve scalable computation and communication [21]. This was the route selected by the Networked Quantum Information Technologies Hub (NQIT) based at the University of Oxford, running from 2014-2019. To perform logical operations on a qubit within a processing node, it must facilitate a strong interaction with external elements. In contrast, to coherently transport a qubit between nodes, it must be inherently non-interacting. As a consequence, a core technical requirement of a quantum network is a coherent interface between qubit hosts, allowing information to pass between stationary processing qubits and travelling qubits.

Our stationary qubits of choice are single atoms, as they may be readily controlled by laser illumination and, within a controlled environment, demonstrate identical energetic structures. Photons are then a natural means of communication, as they can travel significant distances within optical fibres at room temperature, without decoherence. To establish a coherent interface between an atom and single photon, allowing for the transfer of their quantum states, they may be co-confined within a high finesse cavity [22].

1.2 Research Objectives

The objective of this thesis is to develop, explore, characterise and implement the best possible optical cavity for the coherent control of light and matter at the single particle level, fully compatible with modern atom traps. This comprises the competitive benchmarking of state of the art mirror manufacturing techniques, followed by their substantial revision and improvement. With a view towards the realisation of distributed quantum computation, we develop this cavity system for future integration

into a demonstrative quantum network. This will consist of two discrete atom-cavity nodes, connected by photonic communication channels. To establish atomic entanglement, we envisage a similar scheme to that realised with trapped ions [23, 24].

In these foundational quantum networks, heralded entanglement is distributed between the ions by interference and detection of their generated photons. However, the process of spontaneous emission is inherently isotropic and photon collection efficiency is limited by the numerical aperture of the lenses used; in the order of $\ll 10\%$ [25]. An optical cavity acts as a potential solution to this, improving collection efficiency by the process of Purcell enhancement. In moving to the regime of strong atom-cavity coupling, one may realise photon generation schemes with near unit efficiency [26].

The utility of the strong atom-cavity coupling regime is not simply limited to that of enhanced photon collection, but creates an environment in which many quantum information processing protocols may be engineered. This includes cavity-mediated quantum gates [27] or, given the reversibility of the photon generation process, a memory for the storage and retrieval of quantum information [28]. While strong-coupling has been recently achieved in an ion-cavity system [29], the integration of dielectric mirrors with electromagnetic trapping potentials remains a significant challenge [30].

The latter is of no issue when using uncharged particles. Therefore, most effects and phenomena of strongly coupled cavity quantum electrodynamics have been explored using neutral atoms [19]. This has created well-established techniques, but a key challenge in realising an effective quantum interface remains persistent localisation of the atom. While ion-cavity systems are readily able to achieve continuous intracavity trapping lifetimes of several hours [31], the state of the art with neutral atoms is considerably shorter, being on the timescale of seconds [32].

1.2.1 Technical Challenges and Advances in Cavity Fabrication

The fabrication of an optical cavity used for coherent quantum control is a considerable engineering challenge. A high rate of coherent coupling between atom and photon requires a well confined mode volume. The quality of the mirror surfaces must be excellent, to ensure a photon persists within the cavity for an extended interaction period. The length of the cavity must be highly stabilised as to be continually resonant with the desired optical frequency. Photons generated within the cavity must be readily captured, via transmission from the cavity mirrors into a well defined spatial mode. Finally, the atom must be well localised on a cavity antinode, setting the maximum lifetime of the interface.

The first cavities to demonstrate strong coupling to a single atom in the optical regime were formed from super-polished mirrors with dielectric coatings [33]. The reflectivity of mirrors produced in this era has not yet been surpassed [34, 35] and since they were limited to a shallow curvature, an increase in atomic coupling required a reduction to cavity length. This was ultimately limited by penetration of the cavity mode into the dielectric stack [35]. It also comes with a significant reduction to orthogonal access to the cavity mode and thus the ability to insert and localise atomic emitters within it.

The requirement to simultaneously increase atom-cavity coupling and intra-cavity optical access led to a migration from macroscopic super-polished mirrors to micron scale mirrors, or micro-mirrors. These are currently produced using two techniques, laser ablation [36] and focused ion beam milling [37]. However, migrating to these novel glass shaping technologies required significant advances in the overall design and construction of optical cavities. It is important to consider that the physical design of a cavity and the interaction it creates with an atom are fundamentally linked. Thus, for an effective quantum interface, all cavity properties must be carefully tailored.

A frequent challenge is birefringent cavity behaviour, which may impact the generation of pure photon polarisation states and limit their mutual coherence [38]. This is particularly prevalent with mirrors produced by laser ablation, due to inherent ellipticity of the glass shaping laser [39]. Methods have been developed both for the minimisation of ellipticity [40], and for its direct control [41]. However, as well as being elliptical, ablated mirrors also have a Gaussian profile. This leads to higher order mode mixing and the potential for greater optical diffraction loss [42], but also a mechanism to realise stronger atomic coupling [43].

A common practice within optical CQED is to create laser ablated mirrors directly on the tip of an optical fibre, forming a paired fibre-tip cavity [44]. The expected channel of communication between discrete cavity nodes in a quantum network is indeed an optical fibre, as it facilitates the propagation of photons over large distances while protecting them from decoherence. Since photons generated within this cavity may be intrinsically fibre coupled, they constitute a highly scalable platform. However, capture efficiency has been limited by the inherent mode mismatch of the cavity mode with that of a single-mode fibrecore [45]. Recent advances have seen the use of composite fibres with integrated mode matching or photonic crystal fibres [46, 47].

The inherently low mechanical stability of these fibre-tip cavities has proven detrimental to their experimental operation. To sustain resonance with a targeted atomic transition, the length of the cavity must be precisely defined and highly stabilised. Recent work has concentrated on increasing the passive stability of optical cavities, by careful elimination of structural resonances in its design and mounting. Simultaneously, the bandwidth of feedback that may be applied to cavity length stabilisation has been enhanced [48, 49]. The control of photothermal effects in fibre-tip cavity stabilisation has been studied, where heat absorbed from light incident on the cavity generates thermal expansion [50].

1.2.2 Intracavity Neutral Atom Trapping

A coherent light-matter interface in a quantum network requires an optical cavity to engineer a sustained and stable interaction between single atom and photon [21]. Accordingly, the single atom should remain highly localised to a position of maximum field strength within the optical cavity mode. For neutral atoms, the prevailing trapping method is a far-off resonance dipole trap (FORT), generated by red-detuned high intensity laser radiation [51]. However, cavity mediated interactions with single emitters generally require demanding conditions of optical resonance, an enhanced photonic lifetime and a confined mode volume, which are routinely obstructive to dipole trapping light [19]. Therefore, the formation of an intra-cavity single atom trap continues to represent a major objective and driving factor in open access cavity design.

It should be noted that many fundamental exhibitions of coherent quantum behaviour in strongly coupled atom-cavity systems have been performed without intra-cavity atomic trapping [52–54]. In these experiments, a cold and diffuse atomic stream proves adequate, which traverses the cavity mode with a sufficiently slow velocity as to observe the process under study [55]. However, such a scheme is incompatible when engineering interactions between many atom-cavity systems, given the relatively low probability of atomic occupation at any given time [56]. An early improvement to atom-cavity interaction periods was achieved with a standing wave FORT, aligned along the cavity axis, which increased interactions from the order of microseconds to milliseconds [57]. However, a key limitation to this approach is that the dipole trapping light must meet the Fabry-Pérot conditions of optical resonance, which are strongly tailored towards atom-cavity interactions instead. Generally, this leads to poor overlap of the atom-resonant cavity mode with the trapping field, as well as an upper bound to dipole trap intensity and detuning.

Modern variations of this approach make significant use of a free space standing wave

FORT, running orthogonal to the cavity axis [58, 59]. Without the restrictions imposed by cavity resonance, deeper trapping potentials may be established with less induced atomic scattering. In one implementation, a high degree of overlap with the standing wave of the cavity is achieved by applying a variable phase to a retroflected trapping beam [60]. In another approach, the standing wave trap is generated from two beams with independent phase modulation, creating an optical conveyor belt which can transport cold atoms over several millimetres [61]. However, in each of these examples, confinement of the atoms in the axis of the cavity is assisted by the low power and modestly detuned laser used for its length stabilisation. The weaker localisation provided in this axis allows atoms to explore a notable range of the cavity mode, leading to systematic fluctuations in atom-cavity coupling strength.

A recent advance is the fabrication of an optical cavity which supports a standing wave dipole trap in the cavity axis at approximately twice the wavelength of atomic interaction [41]. This was achieved via the fabrication of complex dielectric mirrors with two reflectivity bands, whose phase on reflection was carefully engineered to ensure a high degree of overlap between the two standing waves [62]. This potentially allows for many atoms to be trapped within the cavity, each isolated at a distinct cavity antinode. However, true freedom in the placement of individual atoms within the cavity requires the dipole trap to be fully independent of the conditions of optical resonance [63]. An example is the localisation of two atoms to separate cavity antinodes of a nanophotonic cavity, using a pair of independently steerable trapping beams [64].

In our research, we have developed a cavity independent dipole trap, that facilitates many trapped atoms with flexible positioning. This is achieved by spatial light modulation of the trapping laser, which applies holographic techniques in the creation of a bespoke optical potential landscape [65]. Outside of a cavity, this scheme has been demonstrated to trap arrays of 16 atoms over a period of 10 seconds, with the ability for single atom transport over distances of 20 microns [66]. However, to ensure single

occupation of each atom trap, their spatial profile must be focused using a high-NA lens system, to a waist of approximately 1μm. Future integration with an optical cavity therefore demands a high degree of orthogonal optical access to the cavity mode, as to ensure the converging trapping light does not scatter from the mirrors. However, this conflicts with the basic role of an optical cavity, in establishing a highly confined photonic environment around the atom. A detailed evaluation of the engineering challenges and solutions in building an open-access cavity forms a key motivation of this thesis.

1.3 Thesis Outline

Chapter 2 provides an introduction to cavity quantum electrodynamics, the theoretical framework which describes the interaction of a single atom and photon within a cavity. We consider the formal requirements for a cavity to engineer coherent atomic control and the standard figures of merit used in its description. The process of single photon production is outlined for two and three level atoms.

Chapter 3 concerns the theories of classical optics required for interpreting and tailoring the spectral properties of a Fabry-Pérot resonator. We consider the transverse mode structure that arises from the use of spherical mirrors and its role in cavity characterisation. We describe the operation of dielectric mirrors, elements of their optimal fabrication and their greater role in delimiting the cavity mode volume.

Chapter 4 concerns the design of an optical cavity for coherent quantum control. In examining our specific experimental requirements, we establish a set of geometric, spectral and manufacturing constraints placed on the cavity. By the process of numerical optimisation we evaluate previous cavity implementations and select the mirror fabrication method used in this work.

Chapter 5 outlines our developed apparatus for the laser ablation of mirrors. This combines a platform for automated micro-patterned ablation with one of in situ topographical surface reconstruction. We detail the fabrication of arrays of spherical mirror features and pyramidal tapered substrates. An evaluation of their surface quality is performed and a high reflectivity dielectric coating is applied.

Chapter 6 presents the spectral analysis of our fabricated pyramidal mirrors. Using length scanning spectroscopy, the geometric properties and dielectric coating of the mirrors are evaluated. An extension to mirror characterisation is considered, via the development of a compact ringdown spectrometer. We then give a novel exhibition of multiple resonators realised with just a single tapered mirror substrate.

Chapter 7 details the fabrication and implementation of a tapered science cavity, suitable for the coherent control of single neutral atoms. We outline a novel process used for its alignment and the development of the greater experimental architecture used in its operation. We discuss the integration of atomic cooling and trapping methods, with a view towards establishing an elementary quantum network.

Chapter 8 provides a summary, conclusion and outlook to the work of this thesis.

Chapter 2

Cavity Quantum Electrodynamics

This chapter outlines the theoretical framework of cavity quantum electrodynamics, concentrating on the elements necessary to generate single photons with high efficiency. An understanding of this theory proves critical when discussing the design and fabrication of optical cavities for use in coherent quantum control. We start with a description of the Purcell effect; a fundamental phenomenon pertinent to all atom-cavity systems. By considering the requirements of an effective quantum interface, we then move to the strong coupling regime, illustrating its experimental hallmarks for a two-level atom. By expanding this model to a three level Λ system, now incorporating a driving laser, an optimal method for deterministic single photon production is ultimately presented.

2.1 Introduction to CQED

Our surrounding environment is made visible by the atomic process of spontaneous emission. However, despite its omnipresence, the physical mechanism for this phenomenon remained unexplained for much of scientific history. Even Einstein's treatment of spontaneous emission regarded it to be an immutable property of matter [67]. A quantum treatment, however, outlines that spontaneous emission rates are not fixed, but rather a consequence of an atom's environment [68].

A modern description of spontaneous emission considers it to be akin to stimulated emission, where atomic de-excitation is instead triggered by interaction with a field mode of the vacuum. The existence of vacuum modes comes as a consequence of the zero-point energy of an electromagnetic field. In free-space there exists a continuum

of such modes, giving all coupled atomic energy levels the potential to spontaneously emit. However, by placing an atom within an optical cavity, we constrain the allowed modes, thereby controlling its spontaneous emission properties. In principle, this allows the rate of emission to be greatly enhanced [69], or even inhibited entirely [70].

In cavities with highly confined modes, the strength of the atom-cavity interaction becomes sufficient to display a behaviour similar to a laser driven transition. This allows coherent dynamics between a single atom and cavity photon to take place, the theory for which is well described by the Jaynes-Cummings model [71]. As a consequence, an optical cavity can mediate Rabi oscillations in a two level atom [72], or when combined with a driving laser, three-level atomic population transfer [26]. The latter is of great importance for a range of quantum technological applications, as it allows a single photon to be deterministically emitted into a well defined spatial mode.

2.2 The Purcell Effect

For an initial exploration of the quantum dynamics facilitated by an optical cavity, we need to consider two basic properties. The first is the condition for optical resonance, derived using classical optics [73, 74]. The second is decidedly more quantum mechanical in nature: a determination of the vacuum field strength within a cavity [75, 76]. The following derivations follow a perturbative approach, initially assuming a weak interaction between atom and cavity. However, as will later be demonstrated via the Jaynes-Cummings model, the physical effects elucidated herein are of direct consequence to the strongly coupled systems considered in this thesis.

2.2.1 Optical Resonance

We primarily consider a resonator of Fabry-Pérot design; formed by a pair of facing mirrors [77]. Light confined within this volume will circulate, continuously self-interfere and generate standing waves. The eigenmodes of the resonator are radiation patterns which

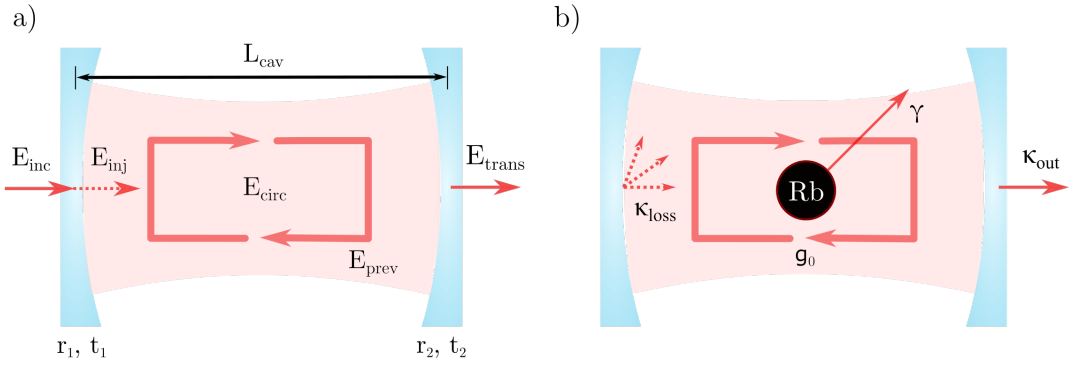

Figure 2.1: **Atom-Cavity Systems. a)** An optical model for a cavity is presented. E_{inc} represents a laser light field incident on one mirror, of which E_{inj} is transmitted into the cavity mode. The steady-state field inside the cavity is given by E_{circ}, which at any given time consists of E_{inj} and the field remaining from a previous round-trip, E_{prev}. **b)** When an atom is present in the cavity (here represented Rb), there are three key rates to consider: atom cavity coupling, g, atomic polarisation decay, γ, and photonic field decay, κ. Note that photonic decay is separated here into propagating losses, κ_{out}, and absorption/scattering losses, κ_{loss}.

are maintained upon reflection. Resonant *modes* are amplified by the cavity, whilst off-resonant frequencies are suppressed. We can derive the conditions of resonance by considering a cavity whose internal photonic field, E_{circ}, comprises contributions from injected laser light and the remaining field from a previous round trip:

$$E_{circ} = E_{inj} + E_{prev}, \tag{2.1}$$

$$E_{circ} = E_{inj} + r_1 r_2 e^{2i\phi} E_{circ}, \tag{2.2}$$

where we assume a steady state field inside the cavity, mirror field reflectivity, r_i, and a round trip cavity phase of 2ϕ. This model is shown graphically in Figure 2.1, identifying the location of the mirrors and fields. In reality, the injected light field is a combination of the transmission coefficient of the input mirror, the laser amplitude incident upon it and the $\frac{\pi}{2}$ phase imparted by propagation through it:

$$E_{circ} = it_1 E_{inc} + r_1 r_2 e^{2i\phi} E_{circ}, \tag{2.3}$$

$$\implies \frac{E_{circ}}{E_{inc}} = \frac{it_1}{1 - r_1 r_2 e^{2i\phi}}. \tag{2.4}$$

A useful quantity to consider is the enhancement to input intensity at a cavity antinode. This is clearly the modulus square of the above expression, normally described as the 'Airy enhancement factor' [78]. For lossless mirrors:

$$A'_{airy} = \frac{I_{circ}}{I_{inc}} = \frac{1 - R_1}{\left(1 - \sqrt{R_1 R_2}\right)^2 + 4\sqrt{R_1 R_2}\sin^2(\phi)}, \tag{2.5}$$

A more comprehensive discussion on this distribution is given in Chapter 3, noting here only the introduction of mirror reflectivities, $R_i = |r_i|^2$. The conditions of resonance are naturally those of maximal enhancement, occurring when the round trip phase is an integer multiple of π. Since the round trip phase is twice the optical length of the cavity, $2L_{cav}$, resonant angular frequencies are given by:

$$\omega_{res} = \frac{\pi c}{L_{cav}} m,\ m \in \mathbb{Z}^+. \tag{2.6}$$

The form of Equation 2.5 is named the 'Airy distribution', shown graphically in Figure 3.1. We consider atomic interactions to be resonant with a single, well defined, cavity mode. The frequency dependence of this mode is found by expanding Equation 2.5 around that resonance, giving a Lorentzian lineshape:

$$I_{cav}(\omega) = \frac{2}{\pi \Delta\omega_c} \frac{\Delta\omega_c^2}{4(\omega - \omega_c)^2 + \Delta\omega_c^2}. \tag{2.7}$$

For a given resonance frequency, ω_c, this profile is clearly characterised by its full-width-at-half-maximum, or cavity linewidth, $\Delta\omega_c$. A determination of the optimal linewidth

is a key element in the cavity design process and will be discussed in detail in Chapter 4.

2.2.2 Vacuum Field Strength

In order to derive the dynamics of a quantised emitter within an optical cavity, one must treat the electromagnetic field quantum mechanically. Given the mathematical complexity of this operation [75], we present only one concise result: this is to consider a cavity as a simple quantum harmonic oscillator [76], whose energy ladder is defined by photon occupation number, n:

$$\varepsilon_n = \left(n + \frac{1}{2}\right)\hbar\omega_c. \tag{2.8}$$

An immediate observation of this spectrum is the existence of a '0-point' energy, $\varepsilon_0 = \frac{1}{2}\hbar\omega_c$, corresponding to a field mode with no photons, also called a vacuum mode. This energy is equally distributed across time averaged electric, $\mathbf{E}_{vac}$, and magnetic fields. Since the atom only couples meaningfully to the electric field, we consider:

$$\varepsilon_0 = 2\int \frac{1}{2}\epsilon_0 |\mathbf{E}_{vac}(\mathbf{r})|^2 d^3\mathbf{r} = \frac{\hbar\omega_c}{2}, \tag{2.9}$$

$$V_m = \frac{1}{|E_0|^2}\int |\mathbf{E}_{vac}(\mathbf{r})|^2 d^3\mathbf{r}, \tag{2.10}$$

where the integral term, V_m represents the geometric mode volume of the optical cavity, normalised by the maximum field amplitude E_0. Combining these equations gives a maximum field of:

$$E_0 = \sqrt{\frac{\hbar\omega_c}{2\epsilon_0 V_m}}. \tag{2.11}$$

To calculate the strength of interaction between this field and a resonant atomic transition, we employ the standard dipole matrix approach. The atom-cavity coupling strength, $g(\mathbf{r})$, is therefore given by:

$$g(\mathbf{r}) = \frac{\boldsymbol{\mu}_{xg} \cdot \mathbf{E}_{vac}(\mathbf{r})}{\hbar}, \tag{2.12}$$

for a dipole moment $\boldsymbol{\mu}_{xg}$, representing a transition between atomic states $|x\rangle$ and $|g\rangle$. We consider the dipole to be aligned to the cavity mode, such that $\boldsymbol{\mu}_{xg} \cdot \mathbf{E}_{vac} = \mu_{xg} E_0$. We also consider the atom to be stationary at an anti-node of the standing wave within the cavity, experiencing a maximal coupling rate g_0:

$$g_0 = \frac{\mu_{12} E_0}{\hbar} = \sqrt{\frac{\mu_{xg}^2 \omega_c}{2\hbar\epsilon_0 V_m}}. \tag{2.13}$$

For an emitter located away from an anti-node, it can be assumed that the coupling strength follows the field, such that $g(\mathbf{r}) = g_0 E_{vac}(\mathbf{r})$. It is clear that the coupling between atom and photon is only strong when the geometric confinement of the cavity mode is significant. Minimising cavity mode volume is therefore often the key to achieving strong coupling, the technical challenges of which are discussed in Chapter 5.

2.2.3 Spontaneous Emission Rates

One route to calculating atomic transition rates is to use Fermi's golden rule, where we normally consider spontaneous emission to be facilitated by a continuum of vacuum modes [75]. For a transition matrix element M_{xg} and density of states $g(\omega)$, the standard rate equation is:

$$W = \frac{2\pi}{\hbar} |M_{12}|^2 g(\omega). \tag{2.14}$$

The derivation of these quantities for free space is given in [76], resulting in a rate:

$$W_{free} = \frac{\mu_{12}^2 \omega_{xg}^3}{3\pi\epsilon_0 \hbar c^3} \equiv \frac{1}{\tau_R}, \tag{2.15}$$

where τ_R is the radiative lifetime of the atom. However, we now anticipate that the vacuum field within an optical cavity does not form a continuum, but rather follows the conditions of optical resonance outlined in Equation 2.11. This is a clear modification

to the density of states local to the atom, coupled with the cavity field enhancement derived in Equation 2.13. This gives alternative quantities:

$$g(\omega_{xg}) = \frac{2}{\pi\Delta\omega_c}\frac{\Delta\omega_c^2}{4(\omega_{xg}-\omega_c)^2+\Delta\omega_c^2}, \tag{2.16}$$

$$M_{xg}^2 = \mu_{xg}^2 E_0^2 = \frac{\mu_{xg}^2\hbar\omega}{2\epsilon_0 V_m} = g_0^2. \tag{2.17}$$

In the following derivation it is convention to introduce a quality factor for the resonator $\mathrm{Q} = \frac{\omega_c}{\Delta\omega_c}$, characterising the sharpness of the resonance.[1] We now repeat Fermi's golden rule to determine the rate of emission into the cavity mode, giving:

$$W_{cav} = \frac{2Q\mu_{xg}^2}{\hbar\epsilon_0 V_m}\frac{\Delta\omega_c^2}{4(\omega_{xg}-\omega_c)^2+\Delta\omega_c^2}. \tag{2.18}$$

In order to quantitatively evaluate the magnitude of this effect we introduce the Purcell factor, F_P. This compares the atomic emission rate within a resonant cavity to the rate outside the cavity:

$$F_p = \frac{W_{cav}}{W_{free}} = \frac{3Q\left(\frac{\lambda}{n}\right)^3}{4\pi V_m}, \tag{2.19}$$

where λ is the wavelength of light and n is the refractive index of the cavity medium. If $F_p > 1$, spontaneous emission is enhanced by the presence of the cavity, whereas for $F_p < 1$, it is suppressed.

The Purcell effect proves valuable for atomic fluorescence collection, as not only does the cavity increase the flux of generated photons, but they are emitted into a well defined spatial mode. A photon detector placed behind a cavity mirror would therefore capture a stronger signal when contrasted against the isotropic nature of spontaneous emission into free space. This has created significant interest for the integration of

[1] The characterisation of a resonator by optical Q factor is uncommon within the close academic community, owing to presence of a similarly descriptive quantity called Finesse. As discussed in Section 3.1, this quantity also characterises resonance sharpness, but is independent of mode number.

optical cavities into trapped ion quantum processors, where quantum state readout is widely performed using high numerical aperture collection lenses, with limited capture efficiency [23, 24]. It will be further demonstrated in Chapter 4 that, with a more comprehensive modelling of the atom-cavity system, one can design an optical cavity which exceeds the Purcell enhancement outlined above [38].

2.3 Strong-Coupling Regime

For coherent quantum control in a light-matter interface, we require system dynamics to be strongly driven by a single photon and atomic excitation [79]. This is not required for the Purcell effect, since a generated photon may escape the cavity before having a sustained effect on the evolution of the atomic state. In the given derivation with Fermi's Golden Rule, the ability of an intra-cavity photon to re-excite the atom is neglected: we have effectively assumed that photons are emitted from the cavity considerably faster than the rate of atom-cavity coupling, g_0.

A more comprehensive model of the atom-cavity system is outlined in Figure 2.1b). Here, the rate of coherent interaction between atom and photon, g_0, is contrasted against the decay rate of atomic polarisation, γ, and photonic field decay rate, κ. The full dynamics of this system are calculated using a quantum master equation, however not without some complexity. By considering the nature of these rates in detail, we shall outline a regime in which coherent atom-cavity dynamics dominate and the idealised Jaynes-Cummings Hamiltonian is a valid representation. This allows a clean exhibition of the physical mechanisms used to build a single photon source that, as with Purcell enhancement, emits photons into a directional mode. However, unlike Purcell enhancement, this photon source will additionally have inherent reversibility [80].

2.3.1 Atomic Polarisation Decay Rate

Within a cavity, incoherent decay of an excited atomic state can take place by two radiative mechanisms. The first is a transition to an uncoupled state, due to which closed atomic transitions are used for quantum information processing. The second is resonant emission into free space, instead of the cavity mode. By solution of the quantum master equation [81], one can determine that the probability of emission into free space from within the cavity, γ, scales with the solid angle subtended by the cavity mirrors, ξ:

$$\gamma = \frac{\Gamma}{2}\left(1 - \frac{\xi}{4\pi}\right), \tag{2.20}$$

where the excited *population* decay rate in free space is Γ. Since cavities considered in this thesis have a small solid angle enclosing the atom, we further assume the free space emission rate to be unchanged, $\gamma = \frac{\Gamma}{2}$.

2.3.2 Photonic Decay Rate

Cavity mirrors are required to have finite transmission, allowing generated photons to be detected, analysed or interfered. We define the rate at which photons are coupled into this propagating modes as $2\kappa_{out}$. However, we are additionally required to consider incoherent decay mechanisms, including scattering or absorption by the mirror surfaces, giving a total rate of $\kappa = \kappa_{loss} + \kappa_{out}$.[2] It can be shown that photonic lifetime is equivalent to the linewidth of the cavity, such that $2\kappa = \Delta\omega_c$ [19].

2.3.3 Cooperativity

Using the nomenclature developed above, we can now state that the Purcell regime was described by $\kappa \gg g_0 \gg \gamma$ [19]. Next we consider the regime of quantum interfacing,

[2] Due to inconsistencies in standard texts, it should be emphasised that we have defined κ in terms of the photonic field, such that $E_{cav}(t) = E_{cav}(0)e^{-\kappa t}$ and $I_{cav}(t) = I_{cav}(0)e^{-2\kappa t}$. Note that a similar factor of 2 was applied to atomic polarisation decay γ, to keep definitions consistent.

specifying the conditions under which single excitations define system dynamics. The photonic mode clearly meets this criteria when it saturates an atomic transition. This can be characterised by atomic emission into the cavity mode exceeding free space, outlining a definition for the critical photon number, n_0:

$$n_0 = \frac{\gamma^2}{2g_0^2}. \tag{2.21}$$

For single photon systems, we naturally require that $n_0 \ll 1$. In analogy, we define a critical atom number, N_0, described as the number of atoms required to meaningfully change the transmission properties of the resonator. For the cavities developed in this thesis, we therefore set $\{n_0, N_0\} \ll 1$ to be our formal requirement. The quantity N_0 was used frequently in the characterisation of optical bistability, by way of its reciprocal value, cooperativity [82]. For the single atom case, we define cooperativity to be:

$$C = N_0^{-1} = \frac{g_0^2}{2\kappa\gamma}. \tag{2.22}$$

This dimensionless constant compares the rate of coherent interaction between atom and cavity to incoherent loss rates. Within the strong coupling regime, it can be shown to be twice the derived Purcell factor, $2C = F_p$ [19]. Finally, it will be demonstrated in Chapter 4 that cooperativity is used to quickly estimate the efficiency of the photon generation process. Given the range of metrics this quantity encapsulates, it is largely considered the standard figure of merit for strongly-coupled atom-cavity systems, for which $C \gg 1$.

2.4 Single Photon Production

In the strong coupling regime, decay mechanisms are minimised, allowing us to consider an idealised system comprising pure quantum states. The time evolution of this system

can then be derived using the Jaynes-Cummings Hamiltonian. A two level atom will be used to demonstrate vacuum Rabi splitting, a phenomena used as an experimental herald of strong coupling [83]. When this is expanded to a three level atom, the mechanism for single photon production can be outlined.

2.4.1 Two Level Atom

We consider a two level atom consisting of a ground state $|g\rangle$, and excited state $|x\rangle$, with energies $\hbar\omega_g$ and $\hbar\omega_x$ respectively. The frequency of the atomic transition is $\omega_{xg} = \omega_x - \omega_g$, giving an atomic Hamiltonian:

$$\hat{H}_{atom} = \hbar\omega_g |g\rangle \langle g| + \hbar\omega_x |x\rangle \langle x| . \tag{2.23}$$

$$\tag{2.24}$$

The Hamiltonian of the photonic mode is similarly defined as:

$$\hat{H}_{field} = \hbar\omega_c \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right), \tag{2.25}$$

where $\hat{a}^\dagger$ and $\hat{a}$ are bosonic creation and annihilation operators respectively. Within the cavity there is a closed exchange of excitation quanta, where absorption of the photon entails excitation of the atom and vice versa. Moving into a rotating frame of reference and making the secular approximation gives an interaction Hamiltonian:

$$\hat{H}_{int} = \hbar g_0 \left(\hat{a}^\dagger \hat{\sigma}_{gx} + \hat{\sigma}^\dagger_{gx} \hat{a} \right), \tag{2.26}$$

where we have defined atomic lowering and raising operators as $\hat{\sigma}_{gx} = |g\rangle \langle x|$ and $\hat{\sigma}^\dagger_{gx} = |x\rangle \langle g|$ respectively. The Jaynes-Cummings Hamiltonian is the sum of these three terms:

$$\hat{H}_{JC} = \hat{H}_{field} + \hat{H}_{atom} + \hat{H}_{int}. \tag{2.27}$$

A common method of treating this Hamiltonian is to split it into two commuting parts [22]. This allows determination of the following eigenfrequencies for the system:

$$\omega_n^{\pm} = \omega_c\left(n + \frac{1}{2}\right) + \frac{1}{2}\left(\Delta_c \pm \sqrt{4ng_0^2 + \Delta_c^2}\right), \tag{2.28}$$

where $\Delta_c = \omega_{xg} - \omega_c$ is the detuning between atom and cavity in a system with n excitation quanta. This energy spectrum is commonly called the Jaynes-Cummings ladder, shown in Figure 2.2a). For a given n, the splitting between dressed states is given by $\Omega_{c,n} = \sqrt{4ng_0^2 + \Delta_c^2}$. In a resonant system with a single excitation, $\Omega_c = 2g_0$. The observation of vacuum Rabi splitting for a single atom is a hallmark demonstration of strong-coupling [83] and remains a key experimental method to determine g_0. Furthermore, the probability of the atom being in the excited state in this system is $\mathrm{P}_{|x\rangle}(t) = \cos^2\left(\frac{\Omega_c t}{2}\right)$, giving Ω_c the title *vacuum Rabi frequency*, as it characterises the frequency with which an excitation oscillates between atom and photon.

2.4.2 Three Level Atom & V-STIRAP

We now expand our system to include a second atomic ground state $|u\rangle$, coupled to the excited state $|x\rangle$ by way of a classical laser field. We define a ground state energy of $\hbar\omega_u$, laser frequency of ω_l and detuning $\Delta_l = (\omega_x - \omega_u) - \omega_l$. This sets up a three level Λ system, as shown in Figure 2.2. The interaction Hamiltonian for this system is:

$$\hat{H}_{int} = \hbar\Delta_l |u\rangle\langle u| + \hbar\Delta_c |g\rangle\langle g| - \frac{1}{2}\hbar\Omega(|x\rangle\langle u| + |u\rangle\langle x|) - \hbar g_0\Big(|x\rangle\langle g|\,\hat{a} + \hat{a}^\dagger |g\rangle\langle x|\Big). \tag{2.29}$$

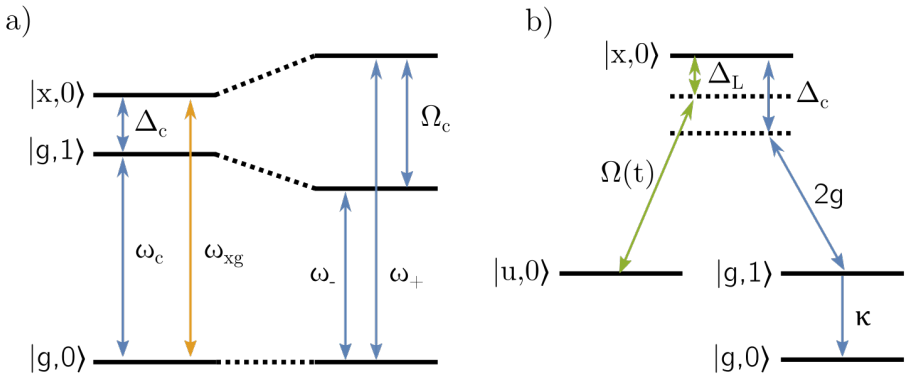

Figure 2.2: **Atomic Energy Structure. a)** The splitting of energy states by action of the Jaynes-Cummings Hamiltonian is shown for the $n = 1$ case. On the left are energy levels for an atom outside of a cavity, with the Jaynes-Cummings Ladder of dressed states on the right. **b)** A three level Λ system is shown, including an uncoupled state, $|g, 0\rangle$, in which the excitation has left the cavity.

The eigenstates of this Hamiltonian are now dressed state triplets:

$$\left|\phi_n^0\right\rangle = \cos\Theta\,|u, n-1\rangle - \sin\Theta\,|g, n\rangle\,, \tag{2.30}$$

$$\left|\phi_n^+\right\rangle = \cos\Phi\sin\Theta\,|u, n-1\rangle - \sin\Phi\,|x, n-1\rangle + \cos\Phi\cos\Theta\,|g, n\rangle\,, \tag{2.31}$$

$$\left|\phi_n^-\right\rangle = \sin\Phi\sin\Theta\,|u, n-1\rangle + \cos\Phi\,|x, n-1\rangle + \sin\Phi\cos\Theta\,|g, n\rangle\,. \tag{2.32}$$

Which, for $n = 1$ consist of states $\{|u, 0\rangle\,, |x, 0\rangle\,, |g, 1\rangle\}$. The mixing angles Φ and Θ have general form:

$$\tan\Theta = \frac{\Omega}{2g_0\sqrt{n}}, \text{ and } \tan\Phi = \frac{\sqrt{4ng_0^2 + \Omega^2}}{\sqrt{4ng_0^2 + \Omega^2 + \Delta^2} - \Delta}. \tag{2.33}$$

Of chief interest to single photon production is the $\left|\phi_n^0\right\rangle$ 'dark' state, in which atomic population is distributed across only the two ground states. The ratio of this distribution is given by the relative strengths of g_0 and $\Omega(t)$. We consider g_0 to be fixed within the timescale of photon production, the system to be Raman resonant, $\Delta_l = \Delta c \equiv \Delta$,

and to be in an initial state $|u, 0\rangle$. With careful temporal shaping of the applied laser pulse, $\Omega(t)$, population can be adiabatically transferred between the ground states, avoiding transient population of the excited state [56]. When the system reaches the state $|g, 1\rangle$, the photon will decay from the cavity at a rate 2κ, leaving the system in $|g, 0\rangle$ and thereby uncoupled from further evolution.

This process is called V-STIRAP, or Vacuum Stimulated Raman Adiabatic Passage [84]. In contrast to collecting fluorescence from rapidly excited atoms, its utility for single photon production is clear. Chiefly, single photons can be produced on demand, synchronised to the driving pulse, $\Omega(t)$ [80]. Since the atom is never in an exited state, spontaneous emission is heavily suppressed. The process is also inherently reversible, allowing one to map the state of one cavity system to another [85]. Finally, the generation process can reach near unit efficiency under the correct conditions, the derivation of which will prove instrumental in designing cavities for coherent quantum control; the subject of Chapter 4.

V-STIRAP is a well established method to deterministically produce single photons [26, 86]. There exist a number of protocols in which a generated photon may be encoded with quantum information. By careful shaping of the laser driving pulse, photon wavepackets with an arbitrary temporal profile may be created [87]. Such a photon may be composed of a number of distinct peaks, or temporal modes, establishing a time-bin encoded information carrier [80]. Narrowband cavity photons have been applied to quantum integrated photonics, demonstrating a controlled-NOT gate [88] and performing multi-mode interferometry [89]. Photons have been generated in well defined polarisation states [90], allowing for their deterministic routing and the potential creation of a many-channel single photon source [56]. Finally, elementary procedures in quantum networking have been realised with similar systems, including coherent atom-cavity interfaces [91] and cavity mediated entanglement between spatially separated atoms [92].

Chapter 3

Optical Theory of Cavities

This chapter presents the theory of classical optics required for the design, characterisation and operation of optical cavities for use in coherent quantum control. Based on the historical application of optical cavities as instruments of spectral analysis and laser feedback, we open with a discussion on the reinterpretation of traditional cavity parameters within the framework of CQED. We discuss the measurement of these parameters in realistic systems, outlining favourable experimental conditions. Our cavity model is then expanded to consider the mode structure that arises from the use of spherical mirrors. Finally, we consider the nature and operation of high reflectivity dielectric mirror coatings and their subtle influence on cavity properties.

3.1 Airy Distribution & Finesse

In the previous chapter we introduced a simple model of a one-dimensional Fabry-Pérot resonator, whose spectral response was described by a comb function, comprising an infinite set of regularly spaced resonant frequencies [78]. This profile was named the Airy distribution, now shown in Figure 3.1a), and is chiefly characterised by the linewidth of each resonance and their separation.

The full width at half maximum linewidth of a cavity is defined as twice the frequency shift required to reduce the internal enhancement factor to half of its resonant value. This can be readily derived from Equation 2.5, by expansion around a resonance, giving:

$$\Delta\omega_c = \frac{\pi c}{L_{cav}} \frac{1 - \sqrt{R_1 R_2}}{\pi \sqrt[4]{R_1 R_2}}. \tag{3.1}$$

The separation of resonant frequencies is described as the free spectral range[1] of the resonator, $\Delta\omega_{fsr}$:

$$\Delta\omega_{fsr} = \frac{\pi c}{L_{cav}}. \tag{3.2}$$

In principle, these two properties are only measurable if each Lorentzian line shape can be resolved from the underlying Airy distribution. Expressed quantitatively, the free spectral range must exceed the cavity linewidth. We encode this condition in a quantity called Lorentzian finesse, which formally equates to the Taylor criterion of spectral resolution when $\mathcal{F} = 1$.

$$\mathcal{F} = \frac{\Delta\omega_{fsr}}{\Delta\omega_c} = \frac{\pi \sqrt[4]{R_1 R_2}}{1 - \sqrt{R_1 R_2}}. \tag{3.3}$$

For strongly-coupled optical CQED with open access cavities, finesse is generally high, in the order of $\mathcal{F} > 10^5$ [19]. The resolution of individual cavity modes is therefore never of direct concern, with the values of $\Delta\omega_{fsr}$ and $\Delta\omega_c$ being dictated instead by experimental practicality, as performed in Chapter 4. However, despite the original definition of finesse being of little applicability to CQED, it has remained a standard metric for the description and comparison of optical cavities. This owes to the existence of several casual interpretations of its meaning. It is generally used in place of resonator Q factors, to indicate the sharpness of resonance. Additionally, it represents a similar value to the internal enhancement factor of a lossless cavity on resonance, thus representing the relative strength of classical intra-cavity effects.

[1] In the description of laser cavities this quantity is commonly referred to as *axial mode spacing.*

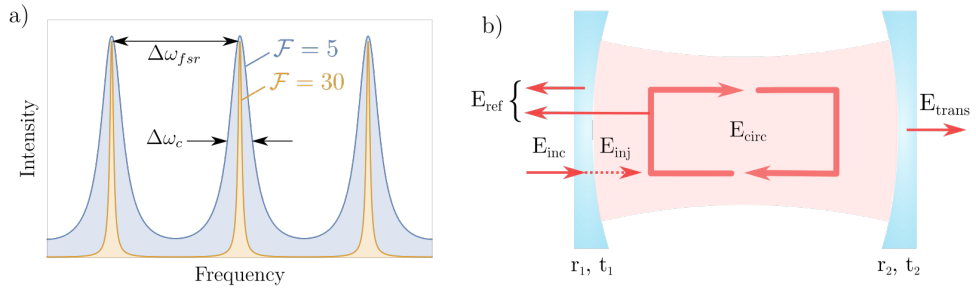

Figure 3.1: **Airy Distribution and Realistic Cavity Model. a)** The Airy distribution consists of an infinite set of regularly spaced resonant frequencies. Each Lorentzian feature is characterised by its linewidth, $\Delta\omega_c$ and each consecutive resonance is separated by the free spectral range, $\Delta\omega_{fsr}$. The 'sharpness' of each resonance is given by the ratio of free spectral range to linewidth, denoted finesse, $\mathcal{F}$. **b)** A realistic model of a cavity considers the origin of the pump laser light used to excite resonance, E_{inc}, of which E_{inj} passes into the cavity. The cavity resonance signal is then measured in two accessible planes: the light transmitted by the cavity, E_{trans}, and the combined back-reflected fields, E_{ref}. Light circulating within the cavity is given by E_{circ}.

3.2 Measured Cavity Signals

The characterisation of an Airy distribution is a key step in deriving cavity properties and tailoring their experimental performance. However, an optical cavity is an intrinsically shielded environment, giving us little ability to measure the intra-cavity field directly. We are therefore required to interpret external signatures of resonance, such as its transmitted and back-reflected fields. To do this, we must make modifications to the current model of optical resonance [73]. Primarily, we must consider the discrete losses that occur at each mirror surface due to imperfections. Then, we propagate the point of interest to outside the cavity.

An updated model for an optical cavity is shown in Figure 3.1b). To consider the source point of the pump laser, we replace E_{inj} in Equation 2.1 by a combination of the transmission coefficient of the input mirror, the laser amplitude incident upon it and the $\frac{\pi}{2}$ phase imparted by propagation through it:

$$E_{circ} = it_1 E_{inc} + r_1 r_2 e^{2i\phi} E_{circ}, \tag{3.4}$$

$$\implies \frac{E_{circ}}{E_{inc}} = \frac{it_1}{1 - r_1 r_2 e^{2i\phi}}. \tag{3.5}$$

In order to introduce optical losses into this model, whilst maintaining concise algebra, delta notation[2] shall be used. This represents reflectivity as $R_i = e^{-\delta_i}$, where for small losses we can Taylor expand to $R_i = 1 - \delta_i$. This implies $\delta_i = T_i + \mathcal{L}_i$, the fractional intensity reduction due to mirror transmission, T_i, and loss, $\mathcal{L}_i$. This gives the field reflectivity coefficient as $r_i = 1 - \frac{1}{2}(T_i + \mathcal{L}_i)$, such that:

$$\frac{E_{circ}}{E_{inc}} = \frac{it_1}{1 - r_1 r_2 e^{2i\phi}} = \frac{i\sqrt{T_1}}{1 - e^{-\frac{T_1+\mathcal{L}_1}{2} - \frac{T_2+\mathcal{L}_2}{2}} e^{2i\phi}}. \tag{3.6}$$

The resonance line shape will naturally be a Lorentzian function, with modifications made to its peak value and linewidth. The expression for linewidth can simply be copied from Equation 2.7, with the peak given by:

$$\left.\frac{I_{circ}}{I_{inc}}\right|_{\omega_{res}} = \frac{4T_1}{(T_1 + T_2 + \mathcal{L}_1 + \mathcal{L}_2)^2}. \tag{3.7}$$

The transmitted field will have a similar form, further including a half-cavity propagation and transmission, such that:

$$\frac{E_{trans}}{E_{inc}} = \frac{it_1 t_2 e^{i\phi}}{1 - r_1 r_2 e^{2i\phi}} \text{ and } \left.\frac{I_{trans}}{I_{inc}}\right|_{\omega_{res}} = \frac{4T_1 T_2}{(T_1 + T_2 + \mathcal{L}_1 + \mathcal{L}_2)^2}. \tag{3.8}$$

Determining the form of the back-reflected signal is slightly more complex. It is cre-

[2] Delta notation dates from the early development of lasers. At this time, optical gains were normally low and therefore high reflectivity cavity mirrors were required for continuous lasing. Presented here is a simplified form of this treatment, with a complete discussion being given by Siegman [73]

ated by interference between the field directly reflected from the mirror and the field transmitted after cavity circulation:

$$E_{ref} = r_1 E_{inc} + i t_1 r_2 E_{circ} \approx r_1 E_{inc} - \frac{r_2 t_1^2 e^{2i\phi}}{1 - r_1 r_2 e^{2i\phi}} E_{inc}. \tag{3.9}$$

Since these fields partially cancel each other out, the Lorentzian line shape is inverted, giving a signal minima at resonance:

$$\left.\frac{I_{ref}}{I_{inc}}\right|_{\omega_{res}} \approx \left(\frac{T_2 + \mathcal{L}_1 + \mathcal{L}_2 - T_1}{T_1 + T_2 + \mathcal{L}_1 + \mathcal{L}_2}\right)^2. \tag{3.10}$$

The spectral profile of the internal, transmitted and back-reflected fields are now fully defined. Their behaviour is split into three categories, as defined by the balance of transmission and loss between the mirrors. If $T_1 = T_2 + \mathcal{L}_1 + \mathcal{L}_2$, there is complete cancellation of the back-reflected signal and the optical power is distributed between losses and the transmitted signal. This is called an *impedance-matched* cavity and provides the ideal configuration for the precise measurement of its optical properties. If $T_1 \gg T_2 + \mathcal{L}_1 + \mathcal{L}_2$, the cavity is described as being *over-coupled*, giving a maximum internal enhancement, at the expense of measurable signal power. The final class is an *under-coupled* cavity, defined as $T_1 \ll T_2 + \mathcal{L}_1 + \mathcal{L}_2$. In this configuration the internal enhancement and external signals are suppressed. However, it is outlined in Chapter 4 why this is commonly employed for atom-cavity systems. An illustration of each of these regimes is shown as Figure 3.2.

3.3 Gaussian Modes

In order to improve the realism of the optical resonator model, it should be confined to practical geometries. Whilst the use of planar mirrors proved perfectly sufficient for the description of spectral line shapes, in reality this would be an exceptionally challenging

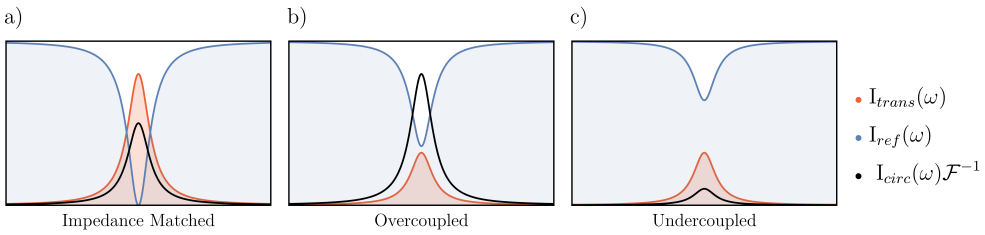

Figure 3.2: **Resonance Lineshapes with Impedance.** The spectral intensity of a cavity scanned across resonance is shown, considered at three planes of reference. The transmitted intensity is shown in orange, the back-reflected intensity in blue and the scaled internal intensity in black: to allow internal intensity to be shown on the same axes, it is divided by cavity finesse. All cavities have equal finite mirror losses and total mirror transmission. **a)** An impedance matched cavity will generate the strongest measurable signals, allowing cavity properties to be measured with high precision. **b)** The over-coupled cavity generates the strongest internal enhancement. **c)** The under-coupled cavity has suppression of all three signals.

device to construct. If the mirrors are not perfectly aligned, the circulating mode will slowly migrate across their surfaces, eventually being lost from the cavity. To counteract this, spherical mirrors are used in CQED, which apply a restorative trajectory to a misaligned beam. Additionally, the use of curved mirrors will focus the cavity mode, increasing the potential atom-cavity coupling strength.

The shape of the cavity mirrors gives the boundary conditions to Maxwell's wave equations, establishing the spatial wavefunctions which can form a stable resonance. In this section we will first consider the fundamental mode of resonance, the cavity geometry required to support it and its effective mode volume. Then, we shall describe the nature of the higher order spatial modes.

3.3.1 Fundamental Gaussian Mode

We define a spatial wavefunction as a radiation pattern which is maintained upon reflection, calculated as an eigenfunction of the paraxial wave equation. Since the standing wave within a cavity must have nodes at the points of reflection, their wavefront curvature must match the spherical profile of the mirrors. This can be fulfilled simply by a

focused Gaussian beam, which indeed forms the fundamental mode of resonance [93].

A more comprehensive discussion of Gaussian optics is provided by Siegman [73], with only the key points outlined here. Within a cavity, the mode geometry is conventionally described using resonator 'g parameters', defined by:

$$\mathrm{g}_i = 1 - \frac{L_{cav}}{R_{oc,i}}, \tag{3.11}$$

where $R_{oc,i}$ is the radius of curvature[3] of mirror i. The mode is fully defined by the size of its waist, w_0, and position of the mirrors with respect to it, z_i. A schematic is shown in Figure 3.3a), with expressions given by:

$$z_i = \frac{\mathrm{g}_j(1-\mathrm{g}_i)}{\mathrm{g}_i + \mathrm{g}_j - 2\mathrm{g}_i\mathrm{g}_j} L_{cav} \text{ and } w_0^2 = \frac{L_{cav}\lambda}{\pi}\sqrt{\frac{\mathrm{g}_1\mathrm{g}_2(1-\mathrm{g}_1\mathrm{g}_2)}{(\mathrm{g}_1+\mathrm{g}_2-2\mathrm{g}_1\mathrm{g}_2)^2}}, \tag{3.12}$$

Where j labels the mirror opposite to the one of interest, i. In the previous chapter, we saw that one of the key properties of the cavity is its effective mode volume. This is derived fully in [94], given here as:

$$V_m = \frac{\pi}{4} w_0^2 L_{cav} \tag{3.13}$$

The Gaussian mode proves to be the most useful for atom-cavity interaction, for two key reasons. Firstly, the central antinode has a simple field distribution, facilitating a steady interaction between atom and cavity. Secondly, the photons generated by this interaction will also have Gaussian spatial wavepackets, allowing them to be efficiently routed and manipulated using standard laser optics. Crucially, they can be coupled

[3]Note that the radius of curvature is the reciprocal of 'mirror curvature', which is more commonly used in casual discussion. As such, the phrase *high curvature mirror* actually denotes a low R_{oc}.

into single mode optical fibres with high efficiency, which forms a key requirement of distributed quantum networking.

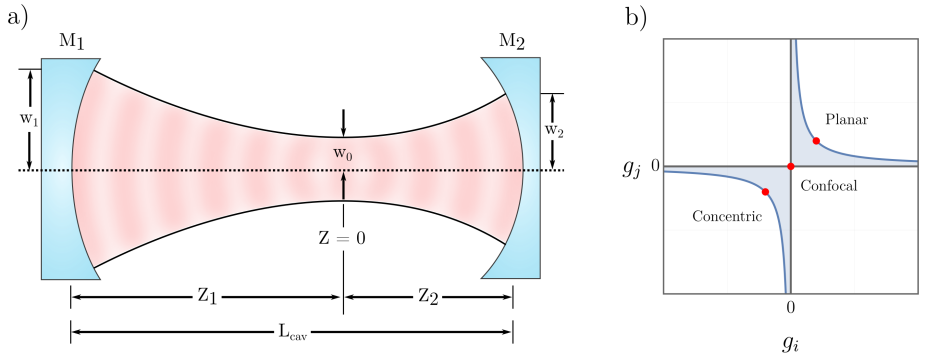

Figure 3.3: **Gaussian Mode Geometry and Stability. a)** The geometry of the fundamental Gaussian mode within a cavity is defined by its spherical mirrors M_i and their separation. Important dimensions include the mode waist w_0, mirror positions, z_i, and mirror spot sizes, w_i. **b)** The shaded region of a stability plot denotes potential resonator geometries which support stable modes of the paraxial wave equation. The position of three hallmark cavity geometries are shown, comprising planar, confocal and concentric mirror arrangements.

3.3.2 Resonator Stability

A consequence of using spherical mirrors is that not all resonator geometries will generate solutions to the paraxial wave equation. Conveniently, the mathematics of this requirement proves to be equivalent to that of a periodic focusing system: for a laser beam to be transmitted over large free-space distances, it must be periodically refocused without significant divergence. Since beam propagation within a cavity consists of repeated reflection from concave mirrors, it can be unfolded and equivalently modelled as a line of regularly spaced lenses. In this model, the focal length of each lens corresponds to mirror curvature. For a given set of mirrors, one can therefore derive a limit to cavity length, after which the beam cannot be effectively refocused and is lost from the system. We define this condition to be the *stability criterion* of the resonator.

The stability criterion can be represented in terms of g parameters as $0 < g_1 g_2 < 1$. It

is convention to represent the range of allowed cavity geometries graphically, by way of the stability plot shown in Figure 3.3b). As it is highlighted in the figure, several points in the stability plot represent geometrical symmetry, associated with interesting optical properties and applications. Two geometries considered in this thesis are confocality, where $R_{oc} = L_{cav}$ and concentricity, where $2R_{oc} = L_{cav}$. Confocal cavities are notable for higher order mode degeneracy and are therefore used for the spectral analysis of light [95]. However, confocal geometries are avoided for single photon generation, as the fundamental mode is desirable for photon capture into a single mode fibre. Concentric cavities represent the maximum stable cavity length and high atomic coupling, but present technical challenges in their fabrication [96]. An equivalent representation to the stability plot is the circle diagram of Deschamps [97], of which a simplified version is given as Figure 3.4.

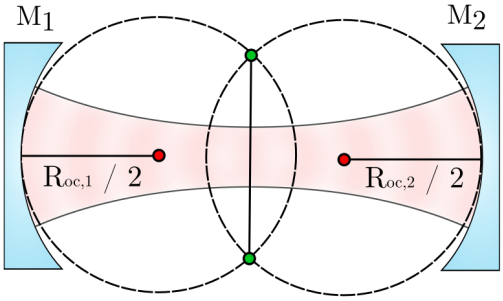

Figure 3.4: **Resonator Circle Diagram.** Each mirror, M_i, is represented as dashed circle of *diameter* equal to its radius of curvature, placed tangent to the mirror surface. The intersection of these circles forms a necessary and sufficient condition for resonator stability. Further, a line drawn between the intersection points (green) outlines the cavity waist position and its relative size. In a confocal cavity, $L_{cav} = R_{oc}$, both circles overlap fully, whereas at concentricity, $L_{cav} = 2R_{oc}$, they overlap only at a single point. If equivalent circles are drawn with a *radius* equal to mirror radius of curvature, then a line drawn between their centres (red) denotes the cavity mode axis. This proves a useful visualisation tool for misaligned cavity mirrors.

3.3.3 Higher-Order Modes

We opened this chapter with a discussion on the *longitudinal* modes within the cavity, later describing them to have a Gaussian intensity profile. However, the fundamental Gaussian mode is only the first element in a complete orthonormal set of solutions to the paraxial wave equation [98]. Each of the higher order solutions are referred to as *transverse* modes of the cavity. The exact shape of these modes is governed by the symmetries of the cavity system in question, with respect to the geometry of the laser beam[4] used to excite it. The dominant symmetry outlines the ideal coordinate system in which to solve the paraxial equation, defining the spatial profile of each eigenmode. Alternatively, one can view this coordinate system as the most natural basis for the decomposition of the excitation beam profile; representing it as a linear combination of spatial cavity eigenmodes. Each potential set of eigenmodes is referred to as a *family* of modes, having specific names for each common coordinate system.

The overlap of the Gaussian laser profile and the cavity mode is defined as the geometric *mode-matching* of the system. For perfect mode-matching, all laser light is naturally coupled into the fundamental Gaussian cavity mode. However, imperfect matching is commonplace, arising from errors in laser waist size, position and the direction of propagation. The exact nature of this error breaks the symmetry of the system, outlining its natural coordinate system. For example, if the propagation axes of the laser and cavity are parallel, but there exists some spatial offset between them, Cartesian coordinates are used. This gives the Hermite-Gaussian family of higher order modes, so named as they are expressed using Hermite polynomials. Alternatively, if the two propagation axes are collinear but there is some offset between waists, the system is represented by cylindrical coordinates. This gives Laguerre-Gaussian modes, expressed using Laguerre

[4]Returning briefly to the situation of vacuum driven CQED, one can immediately note that there is no excitation laser with which to define the symmetry of the system. Instead, the spatial profile of the higher order modes experienced by the atom is defined by the relative orientation and position of the atomic dipole, with respect to the cavity propagation axis. The atom couples to all modes of non-vanishing amplitude, $E_{vac}(\mathbf{r})$, where $\mathbf{r}$ is atomic position.

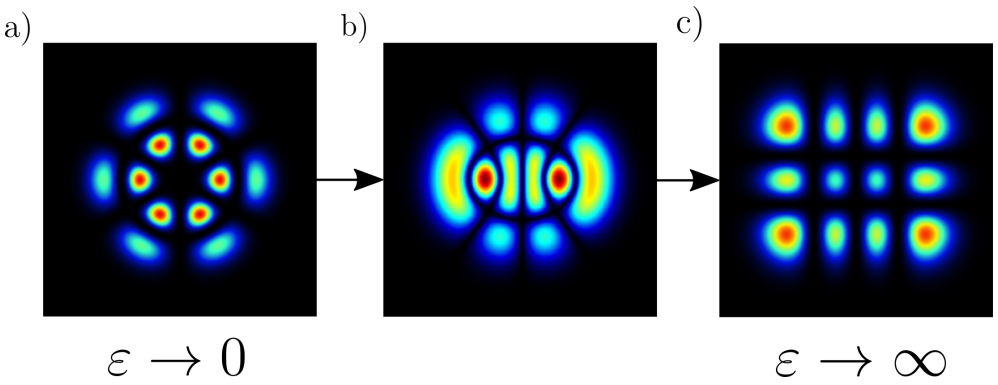

Figure 3.5: **Higher Order Mode Profiles.** The spatial profile of a higher order mode from each of the three core families is shown, as measured in the plane of the cavity axis. **a)** Laguerre-Gaussian mode ($n = 1, m = 3$). **b)** Ince-Gaussian mode ($n = 5, m = 3$). **c)** Hermite-Gaussian mode ($n = 3, m = 2$). However, all modes are indeed Ince-Gaussian modes, with differing eccentricities. At $\varepsilon = 0$, the elliptical coordinate system is equivalent to cylindrical coordinates and at $\varepsilon = \infty$, it is equivalent to Cartesian coordinates. Figure adapted from [99].

polynomials. Example members of each family and their representative geometries are given in Figure 3.5.

In a realistic system there will exist a combination of mode-matching errors. In this general case, the system is solved using elliptical coordinates, yielding the Ince-Gaussian family of modes [100]. This coordinate system is centrally defined by an eccentricity parameter, ε, which at $\varepsilon = \infty$ coincides with Cartesian coordinates and at $\varepsilon = 0$ coincides with cylindrical coordinates. Therefore, in representing the system using Ince-Gaussian modes, we can continuously tune between Hermite and Laguerre solutions.

3.3.4 Transverse Mode Spacing & Gouy Phase

The spectral distribution of higher order modes derives from the Gouy phase: the additional phase acquired by a Gaussian beam as it propagates through a focal region [101]. As this depends on the spatial mode profile, the resonance frequency of each higher order mode becomes non-degenerate. The independent resolution of spatial modes is crucial for single photon production, as we can ensure that the atom is coupled to,

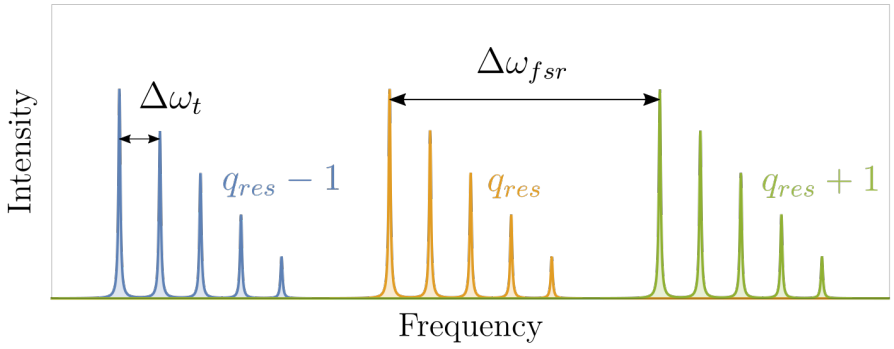

Figure 3.6: **Transverse Mode Spacing.** An example resonance spectra is shown for a near-planar cavity, highlighting the presence of higher order modes for three longitudinal mode numbers, q_i. Each higher order mode is split by $\Delta\omega_t$. Since optical power is distributed across several transverse modes in the example given, some geometric mismatch of the probe beam would be indicated if this was an experimental scenario.

and therefore emits, only into the fundamental mode. Further, this splitting is used in cavity characterisation, as it generates a comb of regularly spaced features used to calibrate the frequency separation between longitudinal resonances. We therefore redefine the resonance condition to be:

$$\omega_{res} = \left(q + (n+m+1)\left(\frac{\arccos\left(\pm\sqrt{g_1 g_2}\right)}{\pi}\right)\right)\frac{\pi c}{L_{cav}}, \tag{3.14}$$

with the $\pm$ term being positive for $\{g_1, g_2\} > 0$ and negative otherwise. Within this thesis, most cavity designs will be in the near-planar regime, defined as $g_1 g_2 \approx 1$. Here, transverse mode frequencies are closely spaced, arranged in consecutive order on the high frequency side of longitudinal resonance. An example spectra is shown in Figure 3.6. The height of each spectral feature indicates the efficiency of geometric mode matching into it. Therefore, if each mode can be simultaneously identified by imaging, a spatial decomposition of the excitation beam can be performed.

As the curvature of the mirrors is increased, the cavity waist becomes further confined,

increasing the extent of the Gouy phase shift. At the point of confocality, even modes[5] become degenerate with longitudinal modes. However, if mirror curvature is increased past this point, transverse mode resolution returns.

3.4 Dielectric Mirror Coatings

The final refinement to our optical cavity model is to discuss the nature of high reflectivity mirror coatings. The development of low loss dielectric mirrors acted as the gateway to the strong atom-cavity coupling regime, with state of the art techniques able to achieve $R > 99.999\%$ across the visible and near-infrared regimes [34]. Dielectric mirrors have replaced metallic mirrors for all high performance applications, whose inherent electron mobility is an upper limit to reflectivity [102].

From a technical perspective, it is important to consider the inherent wavelength dependence of dielectric mirrors, as it affects cavity characterisation and their operation. In the fabrication of mirrors, it is necessary to consider potential incoherent loss mechanisms, as to ensure their minimisation. However, of notable scientific interest is to consider the mechanism by which dielectric mirrors delimit the cavity mode. As will be discussed, this applies a modification to all expressions derived thus far in the thesis and acts to refine the concept of cavity mode volume.

3.4.1 Bragg Stacks

A dielectric mirror, or Bragg stack, consists of many interlaced layers of high and low refractive index materials, deposited on an optically transparent substrate. For a given design wavelength, each of these layers has an optical path length of $\lambda_d/4$, giving rise to its designation as a 'quarter wave stack'. An example schematic is given as Figure 3.7. Since a π phase flip occurs on reflection from the high refractive index surfaces, each of the reflected waves will constructively interfere. Therefore, as monochromatic light

[5] Defined as $m + n = even$.

penetrates further into the stack, its reflection is increased towards unity [103].

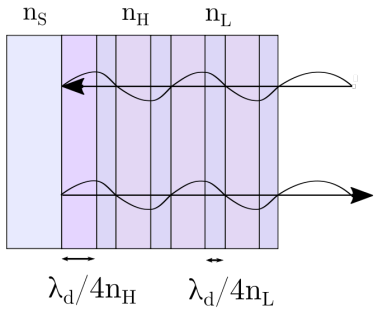

Figure 3.7: **Dielectric Mirror Schematic.** A dielectric mirror consists of quarter wave layers of iterating high and low refractive indices. The light reflected from each interface constructively interferes, thus back-reflecting all incident optical power for a sufficient number of layers. n_H is the refractive index of the high index layer, n_L is the low index layer and n_S is that of the substrate. λ_d is the design wavelength of the dielectric stack, such that each layer has an optical path length of $\lambda_d/4$.

The full dynamics of a dielectric mirror are normally evaluated using scattering matrices. If we assume that the wavelength of light addressing the cavity, λ_{probe}, is equal to the design wavelength of the quarter wave stacks, λ_d, then the reflectivity of the mirror depends only on the number of layers and their refractive indices:

$$R = \frac{(n_H)^{2N} - n_S(n_L)^{2N}}{(n_H)^{2N} + n_S(n_L)^{2N}}. \tag{3.15}$$

Where N is the number of layer pairs and n_S, n_H and n_L are the refractive indices for the substrate, high index layers and low index layers respectively. For 780nm design wavelength, the materials selected are usually tantalum pentoxide, Ta_2O_5, for high index ($n_H \approx 2.28$) and silica, SiO_2, for low index, ($n_L \approx 1.46$). Many layers are deposited for near unit reflectivity, with the first and last layer normally being the high index material. However, a surface layer of Ta_2O_5 has been observed to degrade cavity finesse in vacuum, likely as a result of laser-induced oxygen deprivation [104]. While

this may be reversed by exposure to oxygen and illumination by UV light, it may also be prevented by applying silica as the final layer.

The reflective properties of a dielectric mirror also depend strongly on wavelength. For a simple quarter wave stack, the spectral profile is generally characterised by a broad stop-band around the target wavelength, with oscillating outer lobes. An example transmission curve is given in Figure 5.15b). As the illumination light is detuned from the design wavelength, the effective linewidth of the cavity will increase. Accordingly, the spectral properties of the cavity, as derived above, are generally only valid within a nanometre of design wavelength.

3.4.2 Incoherent Losses

In the fabrication of high quality dielectric mirrors, incoherent loss mechanisms are carefully minimised. Indeed, the reflectivity of a cavity mirror is usually designated by its parts per million (ppm) losses. These arise foremost from the scattering of light by surface imperfections, or absorption by contamination within the coating. Since these losses occur on each round trip of light within the cavity, their effect on the intra-cavity field is analogous to mirror transmission. However, transmission is generally considered a coherent loss mechanism, as it may be utilised in the injection or extraction of light.

Scattering losses are generated by surface imperfections that are on scale with the wavelength of incident light. In standard optical applications, these are usually characterised by the number of scratches and digs over the mirror surface. However, deviations of this magnitude would prevent the formation of a high finesse resonator entirely. Of more routine concern in high finesse resonators is the atomic scale roughness of the surface, which induces losses in the ppm regime. The magnitude of scattering losses can be estimated as [44, 105]:

$$\mathcal{L}_S = \left(\frac{4\pi\sigma_{\text{rms}}}{\lambda_{probe}} \right)^2 \tag{3.16}$$

Where σ_{rms} is the root mean square surface roughness. Absorption losses, $\mathcal{L}_A$, are generally induced by contamination during the coating procedure and are accordingly more difficult to predict. However, they can be minimised by using low-loss coating methods, namely ion beam sputtering (IBS), and post-deposition thermal annealing of the mirrors [106]. The magnitude of scatting losses is not observed to depend on the the applied coating, being purely defined by the base substrate. However, as mirror transmission is lowered by increasing the number of dielectric layers, this gives a greater susceptibility to absorption losses.

3.4.3 Mode Boundary

In our derivation of atom-cavity coupling strength, the cavity mode volume was considered as the geometrical region between the mirror surfaces. This inherently assumed that mirrors constitute a perfect optical boundary to the cavity mode. However, given that light penetrates into the dielectric stack, some component of the cavity volume must be within the mirror itself [107]. Then, since all realistic dielectric mirrors have finite transmission, there must be coupling of these modes to those of the environment.

The influence of dielectric mirrors on atom-cavity coupling is examined in detail in our recent publication [108], for which I was a co-author. This paper considers the cavity as an open-quantum system, where the electric field modes are unconstrained and the presence of the cavity results in a frequency dependent ratio of internal and external field amplitudes. This gives rise to a closed expression for the modified atom-cavity coupling strength, which allows variance of the design wavelength of the dielectric mirror, the resonance frequency of the cavity and the transition frequency of the atom.

The optical path length of the cavity is clearly modified as light penetrates into the dielectric stack. This applies an additional phase to the reflected beam, shifting the conditions of resonance. In literature, the modified spectral properties of the cavity are re-characterised using an effective cavity length, L_{eff} [35]. In general, L_{eff} is not equal to L_{cav}, occurring only when the resonance frequency of the cavity matches exactly the design wavelength of its stack.

Furthermore, when including the dielectric stack in the quantisation of the electromagnetic field, the normalisation volume of the cavity no longer corresponds to L_{cav} or indeed L_{eff}. Instead, it gives rise to an independent effective 'coupling factor' length, $L_N^{(m)}$, where m is the longitudinal mode number and N is the number of high index dielectric layers[6]. $L_N^{(m)}$ only asymptotically approaches the geometric length of the cavity at extended mirror separations. For cavity geometries that are on scale with the wavelength of operation, the modification to the cavity resonance profile and atom-cavity coupling strength are notable. It also raises interesting questions on the formally correct interpretation of cavity mode volume. As is described in the paper, only when the Airy function of the cavity can be represented as a sum of perfectly distinct Lorentzians, does the normalisation factor in atom-cavity coupling strength correspond to a geometric quantity.

However, for the reminder of this thesis, we only consider cavities where $L_{cav} \gg \{\lambda_{probe}, \lambda_d\}$. Therefore, these effects and length variations can be safely ignored.

[6]These quantities are labelled in the same manner as the source publication, for easier reference.

Chapter 4

Cavity Design and Optimisation

This chapter discusses the design considerations of an optical cavity used for coherent quantum control, concentrating on the elements necessary to facilitate strongly-coupled interactions with trapped neutral atoms. First, we consider the geometric and spectral restrictions placed on cavity design by mirror manufacturing methods and experimental suitability. Then, we numerically evaluate the properties of previous optical cavities used in this application, motivating and contextualising the technical developments in cavity fabrication made by this thesis.

4.1 Design Methodology

Our intention is to build an optical cavity system that develops and demonstrates several key aspects of a coherent atom-photon interface. As we have discussed, the foundational requirement for such an interface is to facilitate strong-coupling between single quanta of light and matter. However, this interaction must be sustained, by appropriate localisation of the atom within the cavity volume. This atomic trapping method must not interfere with the operation of the cavity, allowing it to maintain a desired resonance frequency.

An 'ideal' cavity for quantum information processing may only be described if the greater system architecture and computational protocols are perfectly defined. However, given the high rate of development within this field, the route towards an optimal cavity remains subject to alternate viewpoints [21, 109]. In this light, we initially consider broad perspectives on cavity design. As our process becomes more refined, our

specific experimental objectives will be considered in greater detail.

We shall first consider the degrees of freedom available in cavity design and limitations applied by the manufacture of high reflectivity mirrors. Then, we shall consider the geometric and spectral constraints placed on the cavity in order to realise standard experimental objectives. Finally, we shall develop and discuss appropriate metrics of performance for the system, to allow for its numerical optimisation.

4.2 Degrees of Freedom

To simplify the cavity design process, the tailored degrees of freedom must be minimal. As we have seen in Chapter 3, cavity mode geometry is defined by the curvature of each mirror and their separation. However, as will be described in the following chapter, we must now consider the diameter of the mirrored surface, σ_m, and that of the substrate which supports it σ_s. These will play a role in the geometrical limits of an optical cavity, by way of its maximum length and suitability for atom trapping.

To describe the dielectric coating applied to the mirrors, we consider our requirement to efficiently extract and capture photons generated within the cavity. This naturally implies the reduction of all incoherent optical loss mechanisms, with some finite transmission engineered into only one mirror. This ensures the directional output of photons and greatly assists in the experimental simplicity of their capture. We therefore tailor the coating simply via the transmission of this out-coupling, T_{oc} mirror. Transmission through the opposing highly reflecting mirror, T_{hr}, remains useful to classically probe cavity resonance. However, a residual value inherent to the coating process proves sufficient for this.

Despite the large range of parameters used to characterise optical cavities, the degrees of freedom in their design are markedly few. However, the design process is not without

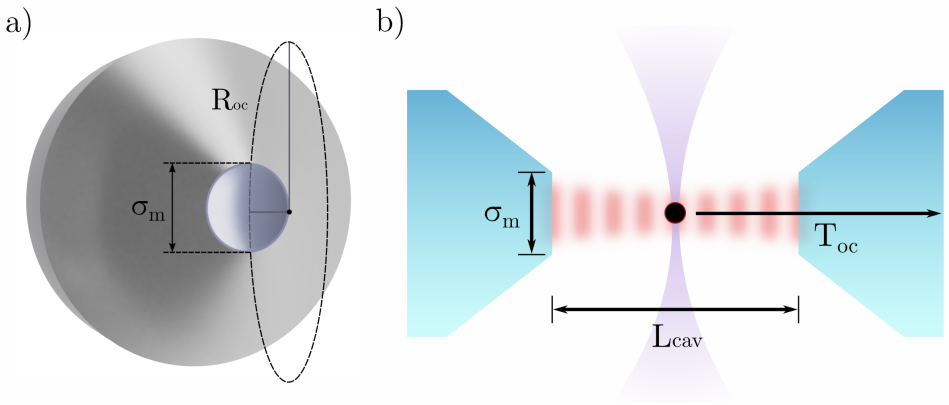

Figure 4.1: **Degrees of Freedom in Cavity Design.** A visual representation is given for the simplified set of degrees of freedom used in our cavity design process. **a)** A single tapered mirror is specified by its radius of curvature R_{oc} and coating diameter, σ_m. In this example, $\sigma_s = \sigma_m$, as the full facet diameter supports the coating. This coating is described by transmission of the output-coupling mirror, T_{oc}. This assumes that incoherent losses and residual transmission of the high-reflecting mirror are automatically minimised. **b)** Cavity length is given by L_{cav}. A basic illustration of an atom-cavity system is given, with the atom in black, the cavity mode in red and dipole trapping beam in purple.

complexity. As each characteristic has an interrelated dependency on the same small set of freedoms, their optimisation is an exercise in trade-off and mitigation. For simplicity, we consider here only geometrically symmetric cavities, where each mirror has identical curvature and diameter. We can therefore describe a cavity design by statement of L_{cav}, R_{oc}, T_{oc}, σ_m and σ_s; a visual representation of which is given as Figure 4.1.

4.3 Manufacturing Limitations

The technical challenges presented in manufacturing high reflectivity dielectric mirrors have been an ongoing limitation to the realisation and scalability of coherent atom-cavity interfaces. Indeed, the landmark demonstration of strong coupling to a single atom in the optical regime has been notably attributed to efforts in reducing dielectric coating losses [34]. This generation of cavity mirrors, comprising super-polished substrates and ion beam sputtered (IBS) coatings, facilitated many subsequent high-

impact studies and was adopted widely across the field [19].

However, given the exceptionally high reflectivity achieved by these mirrors, R > 99.999%, it was not anticipated that system cooperativity could be further improved by reducing optical losses. Instead, the atom-cavity coupling rate would need to be increased via the minimisation of mode volume. This led to the demand to produce higher curvature glass substrates, which retain the optical quality of super-polishing.

This work is ongoing, split between two key methods for glass shaping. These are focused ion beam (FIB) milling [37, 110] and laser ablation [36, 44]. The platform selected has significant implications for the design of optical cavities, in the range of geometries and surface qualities that can be realised. Therefore, we shall now detail our initial analysis of each technique and make a comparison of their regimes of operation.

4.3.1 Super-Polished Mirrors

To create super-polished mirrors, a particulate suspension is driven against a glass surface, shaping it into a spherical profile by the action of abrasion. As progress towards the intended geometry is made, the particulate size is reduced, thereby increasing surface quality. Surfaces approaching atomic smoothness can be achieved, surpassing $\sigma_{rms} = 1$Å. However, this process is limited to the creation of shallow depressions. Indeed, mirrors of high optical quality prove difficult to source commercially for $R_{oc} \ll 20$mm.

To improve access to the cavity mode, it is standard to grind cylindrical super-polished substrates into a conical profile [35]. At the apex of this cone, an untouched portion of the mirror surface should remain, as illustrated in Figure 4.2. However, there is a limit to the final facet size, occurring when the apex shatters or damage to the remaining dielectric surface is substantial. For the mirrors used in Chapter 7, this limit was observed at a facet diameter of $\sigma_s \approx 300\mu$m for a 45° coning angle. We assume this to

be a representative limit to super-polishing for the remainder of this chapter.

4.3.2 Focused Ion Beam Milling

In FIB milling, a focused beam of gallium ions sputters glass from a target substrate, gradually eroding into its surface [111]. The restricted size of this focus[1] allows for nanometer scale structures to be created with high precision. However, given the relatively low flux of ions involved, there is a challenge in expanding the scale of these structures towards standard cavity mirror geometry. To evaluate the use of FIB milling for our cavity production, we produced test structures on glass using a Zeiss Crossbeam 540.

The crossbeam is named for the scanning electron microscope included in the system, for non-destructive sample navigation. While the ion-beam is able to generate an image of the target surface, it will simultaneously erode it. However, since the electron microscope is unable to image a non-conductive surface in the first instance, some metallic material was deposited on the substrate in order to perform alignment. The removal of this material was then imperative to prevent the introduction of contamination into the dielectric coating. Instead of applying a conductive material, the surface may be imaged via secondary ions, but this was not tested in the apparatus available.

Over the duration of milling, Gallium ions gradually become implanted into the mirror site, building a static charge distribution. This increasingly deviated the trajectory of the incoming ion beam, warping the geometry of milled features. To overcome this, an electron flood gun was used, which filled the chamber with a charge screening cloud of electrons. However, the electron microscope naturally becomes over-exposed during this process, preventing visual monitoring of the milling process.

[1] Since this is a non-optical technique using magnetic lenses, the standard diffraction limit experienced in laser ablation does not apply.

The most limiting issue experienced was the maximum volume of material that could be removed. This was a combination of the limited ion current and the stability of the mounting arrangement for long-duration exposures.[2] Further, at the time of research, the surface roughness of milled features was limited to $\sigma_{rms} \approx 1$nm: an order of magnitude greater than desired. This creates notable mirror scattering losses at 780nm, approximately 260ppm, preventing the coherent extraction of photons in our planned experimentation.

The application of FIB milling towards optical cavity production remains an active area of research, particularly within the field of nitrogen vacancy (NV) based quantum technologies [112]. It has been used in the production of large arrays of cavity mirror features on the same substrate, allowing the formation of multiple cavity modes [37]. This is of interest when engineering many emitter-coupled systems, such as coupling to a molecular fluid [113]. This has the potential to enhance the optical interrogation of a chemical sample to yield its constituents, or modify its rates of reaction [114]. The development of laser ablation towards the production of similar cavity arrays is given in Chapter 5.

4.3.3 Laser Ablation

The final technique of discussion is laser ablation, which was initially evaluated using a system developed at the University of Sussex [40, 46].[3] Following this extended visit, a novel implementation was developed within our own laboratory, which is the subject of the following chapter. In order to avoid repetition, we give only an introduction here.

Laser ablation uses a strongly focused laser beam to heat glass, rapidly evaporating near-spherical profiles from its surface. Simultaneously, the underlying molten layer exerts surface tension, pulling the solidifying material into a smooth profile. In utilising a

[2] The standard cost of FIB milling was also a limitation, given the significant overhead required to further develop the process and produce each mirror.

[3] Courtesy of the research group of Professor Matthias Keller.

natural thermal process, one can greatly decrease the engineering overhead required for glass shaping, whilst achieving very high surface quality. However, control of the evaporative and molten process proves difficult, setting tight limits on feature geometry [115].

In Sussex, laser ablation is performed on the facet of optical fibres, which is a common practice within the field of optical CQED. A cavity formed from a pair of fibre-tip mirrors is considered to be inherently scalable, as light can be injected or captured from the mode without discrete optical components [29]. This is of particular value in space-restricted environments, such as ultra-high vacuum, where the requirement for optical beam access to the cavity can be limiting. However, as supported by the work of Chapter 5, laser ablation is practical for a range of complex substrate geometries. In particular, we consider some negative qualities of fibre-tip cavities for forming a quantum interface. These range from challenges in ensuring efficient cavity mode-matching and coherent photon extraction, to low vibrational stability and high photothermal sensitivity. We discuss the development and demonstration of a viable alternative to fibre-tip mirrors, namely pyramidal micro-mirrors.

4.3.4 Comparing Fabricated Geometries

In order to select the mirror fabrication method used in the production of our trapped-atom cavity system, they are compared by the optical cavity properties made accessible. This is assisted by describing each method as a continuous range of possible mirror geometries; defined by their diameter and curvature. However, this treatment is somewhat subjective, overlooking technical specifics of each system in question.

To generate a rough estimation of the geometrical range of each technique, we consider indicative limits on the volume of glass they can remove from a substrate. By consultation of a number of studies that created high quality mirrors [40, 44, 47, 116, 117], the map of Figure 4.2b) has been populated. The quality here is defined as σ_{rms} <0.35nm, corresponding to $\mathcal{L}_S < 32$ppm at 780nm, which for FIB milling required post-thermal

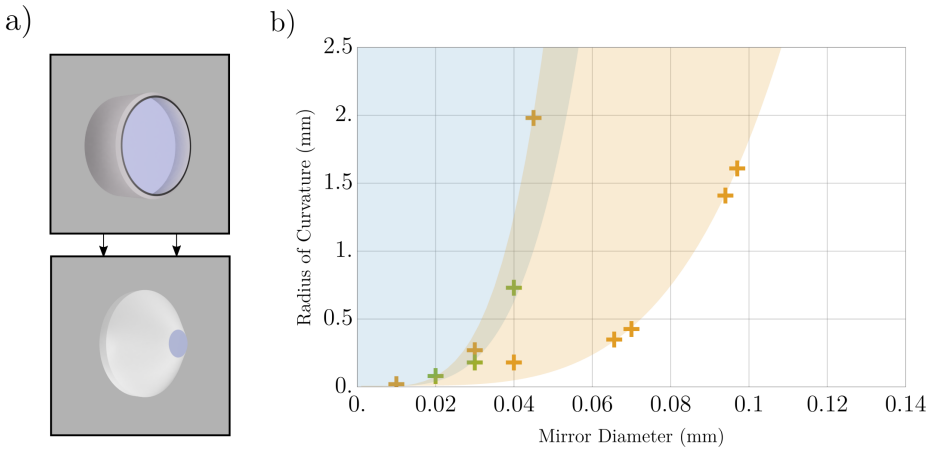

Figure 4.2: **Mirror Fabrication Techniques. a)** A rendering of an unconed mirror is shown in the upper inset, with a diameter of $\sigma_m = 7.75$mm and a coned mirror in the lower, with a reduced diameter of $\sigma_m = 300\mu$m. **b)** Laser ablation and FIB milling are compared by the geometry of spherical features they are able to produce. By collecting limiting values of (σ_m, R_{oc}) reported in a number of studies [40, 44, 47, 116, 117], the region shown has been populated. In order to estimate a continuous range of operation for each technique, we calculate the volume of glass removed from a planar base substrate associated with each coordinate. For laser ablation, coordinates are grouped into maximum and minimum volumes. Lines of constant volume are fitted to these sets, and the region between them shaded in yellow. This region estimates the operational range of laser ablation. An equivalent treatment is given to FIB milling, establishing its operational range in blue. The lower bound of FIB milling is arbitrarily small on the plotted scale.

annealing of the surfaces to smooth inherent roughness.

While this is a highly simplistic model, it represents the collected data appropriately and conveys an overall picture of the state of the art. Generally, FIB milled features are limited to smaller geometries than laser ablation, in line with the physical mechanism of sputtering occurring at the atomic level. The direct route to increasing milled feature depth, via an increased ion current or duration, is not expected to improve surface quality. Therefore, post-thermal annealing will continue to be required to reduce RMS surface roughness to the desirable 0.1-0.3Å range. However, depending on the depth of the molten layer required to smooth surface defects, annealing may alter the

underlying surface profile. The laser system used for thermal annealing is the same as that for ablation, albeit with weaker exposures applied to the glass surface. This secondary manufacturing stage comes with additional complexity and an increased risk of surface contamination. However, the application of laser annealing to milled structures presents as the most immediate route towards the creation of high-finesse cavities using FIB milling in this geometrical regime.

The thermal process underlying laser ablation is decidedly more macroscopic in action. Heat naturally diffuses through a glass substrate, with material evaporation and surface tension occurring on the scale of microns. Since the spread of heat cannot be fully controlled, it can prove difficult to fabricate highly confined geometries. In the first instance, one may attempt to reduce feature diameter by more tightly focusing the ablation laser. However, the focus is at best diffraction limited, which for the standard wavelength range of 9.3μm to 10.6μm is relatively sizeable. One may attempt to reduce feature depth by decreasing exposure power, although this may cause the glass not to reach the critical temperature required for evaporation. This may instead generate an overly thick molten layer, pulling into a convex curvature under surface tension. Overall, if exceptionally small features are required for cavity fabrication, laser ablation may prove a challenging platform for their development.

In the geometrical range of ablation highlighted above, the process comes with the strong advantage of inherently low surface roughness. While laser exposure parameters must be optimised to produce the geometry of choice, the presence of a molten layer is generally assumed. This allows one to create high finesse cavity mirrors with significantly less experimental complexity than FIB milling.[4] Indeed, the current technical challenge comes in further increasing the standard diameter of ablated features. This forms one of the key motivations for the work in Chapter 5, where the topic is discussed

[4]Specifically, we note that FIB milling apparatus is a sizeable and sustained departmental investment, in contrast to the laser system developed in Chapter 5.

in detail. In Chapter 6 we consider the mitigation of issues inherent to ablated mirrors, through potentially elliptical geometries and non-spherical profiles. Our comparison of FIB milling and laser ablation is continued later in this chapter, in the context of our specific experimental requirements.

4.4 Geometric Limitations

We now consider the limitations placed on cavity geometry by our intended experimentation, a core aspect of which is the integration of an atom trapping laser into the cavity volume. This laser propagates orthogonally to the cavity axis, requiring its mode to be openly accessible in this direction. However, simple modifications to cavity geometry to improve access cannot be performed without consequence, as we will now address.

The minimum optical access required of our cavity is given by the focusing optics used in the formation of the dipole trap [65, 66]. This is a complex lens system[5] with an effective focal length of 30mm and numerical aperture of 0.6. The waist of this focus is routinely varied to modify trap properties, by changing the diameter of collimated beam input. To give a general specification of the system, a standard input beam diameter of 20mm creates a 1.6μm waist at the cavity axis. Given the high optical power contained within this focus, partial illumination of the cavity mirrors may generate heating and length instability. Simultaneously, applying this aperture to the beam may deform the potential landscape of the trap, thereby reducing its average lifetime of atomic occupation.

To ensure that the mirrors are suitably distanced from the trapping laser, we demand optical access of $\theta_{oa} = 45°$ with respect to the central cavity antinode. This is a conservative estimation, chosen to be several degrees greater than the numerical aperture of

[5]This lens system was designed to achieve coincident foci at the off-resonant trapping wavelength, 1064nm, and atomic fluorescence wavelength, 780nm, while accommodating for the optical window of the vacuum chamber.

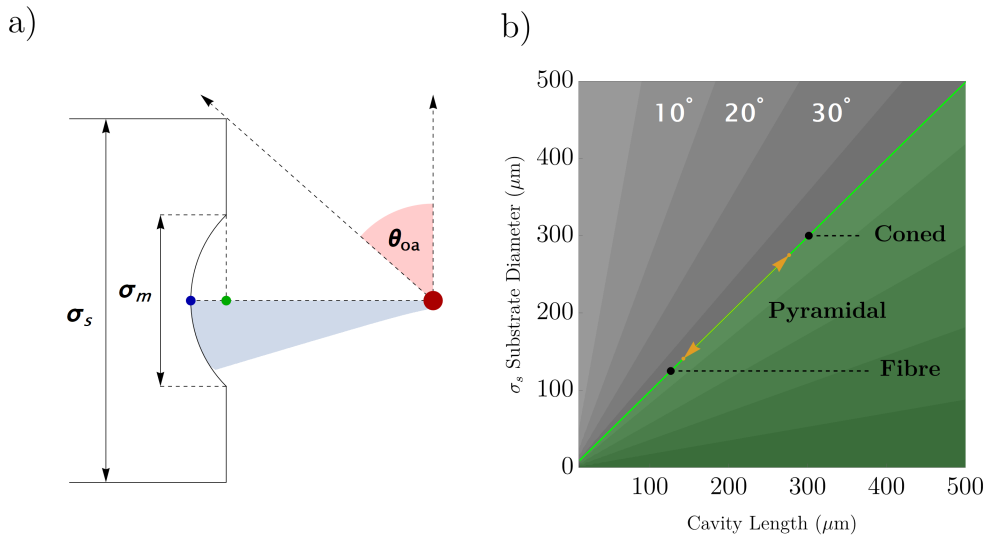

Figure 4.3: **Optical Access. a)** The geometry of optical access is shown, highlighting the distinction between substrate facet diameter and mirror surface diameter. In calculating the distance of the trapping site to the mirror (red to blue circles) one must subtract the sagitta from the cavity length (blue to green circles). **b)** The cavity length requirement for $\theta_{oa} > 45°$ is shown in green, with standard facet diameters highlighted. These include fibre facets, $\sigma_s = 125\mu$m, the pyramidal substrates developed in the following chapter, $\sigma_s = 141 \rightarrow 275\mu$m, and super-coned substrates, $\sigma_s = 300\mu$m.

the lens system. The geometrical definition of optical access is described in Figure 4.3, with an associated mathematical expression of:

$$\theta_{oa}(z_{atom}) = \frac{\pi}{2} - \arctan\left(\frac{\sigma_s}{2\left(\left|z_{atom} - \frac{L_{cav}}{2}\right| - R_{oc} + \sqrt{R_{oc}^2 - \left(\frac{\sigma_m}{2}\right)^2}\right)}\right). \tag{4.1}$$

One might assume from this expression that cavity length is ideally increased towards concentricity, limited only by the stability criterion [118]. However, in many practical scenarios, the length of the cavity is instead limited by mirror surface diameter [44]. Despite the fundamental Gaussian mode of the cavity having a concentrated intensity profile, there is no hard limit to its transverse extent. Therefore, some portion of the radiation field will always exceed the area of the mirror and be lost on reflection. Since

the cavity mode diffracts on propagation, this loss is experienced on each round trip. Accordingly, these 'clipping losses' can be considered similar in effect to those generated by absorption and scattering, albeit with a dependence on cavity mode geometry. A standard model for the total losses generated by each cavity mirror is given by:

$$\mathcal{L}_T(L_{cav}, R_{oc}, \sigma_m) = \mathcal{L}_A + \mathcal{L}_S + \mathcal{L}_C(L_{cav}, R_{oc}, \sigma_m), \tag{4.2}$$

$$\text{where } \mathcal{L}_C(L_{cav}, R_{oc}, \sigma_m) = e^{-\frac{\sigma_m^2}{2w_i\left(L_{cav}, R_{oc,i}\right)^2}}. \tag{4.3}$$

As is observed in Figure 4.4, clipping losses are largely negligible until the size of the mode on each mirror approaches its radius. A large mirror diameter is therefore crucial in preventing clipping losses in long cavities, thereby permitting optical access. However, modifying the extent of clipping losses can also be performed by adapting the mode size on each mirror via their curvature. For a given cavity length, clipping losses are minimised around the point of confocality.

In increasing cavity length to improve optical access, we have required a large mirror surface diameter σ_m. However, optical access equally required a small facet diameter, σ_s. Thus, to achieve the greatest flexibility in cavity design, they are ideally equal. In the standard nomenclature of optics, this is referred to as a full clear aperture of the mirrors. However, this is not readily achieved when using micro-mirrors produced by laser ablation. Our work to increase the clear aperture of ablated mirrors, and the technical issues experienced therein is reviewed in Chapter 5.

A requirement to reduce the extent of the dielectric surface which encloses the atom is not unique to our application, but common to general emitter localisation within optical cavities. A notable comparison is ion-cavity systems, where electromagnetic trapping methods are employed. The presence of dielectric mirror surfaces can distort

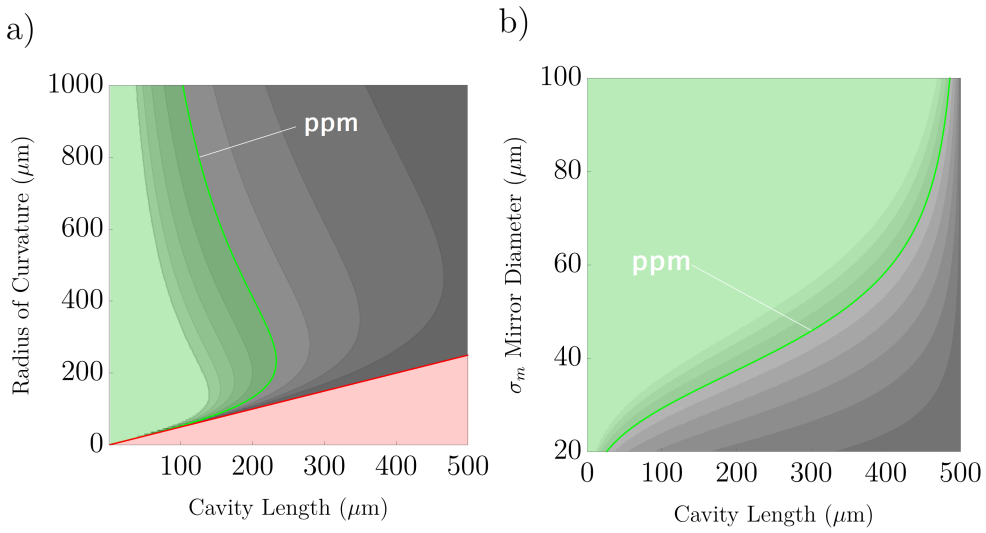

Figure 4.4: **Clipping Losses.** **a)** The relationship between cavity mode geometry and clipping losses is shown for mirrors of $\sigma_m = 40\mu m$. Regions of $\mathcal{L}_c < 1$ppm are shown in green and geometric instability in red. Contours are drawn every $\mathcal{L}_c = 10^x$, where $x \in Z$. For a given length, clipping losses can be minimised by suitable choice of R_{oc}, corresponding to a minimisation of the mode size on the mirror, w_i. **b)** For $R_{oc} = 250\mu$m, the impact of σ_m and L_{cav} on $\mathcal{L}_c$ is shown, using the same colouring as a). For sufficiently large mirror diameters, the onset of clipping losses becomes negligible. However, this is in competition with our abilities in mirror fabrication.

these potentials, requiring mirrors to be located far from the emitter [43, 119]. Thus despite the dissimilarity in the trapping method, their implied restrictions on cavity geometry are similar. Accordingly, development in cavity fabrication in one field is routinely applied to the other.

4.5 Spectral Requirements

4.5.1 Suitability for Length Stabilisation

In the context of this thesis, an optical science cavity is required to provide sustained resonance with a selected atomic transition frequency; specifying a precise set of its allowed length values. In the following chapters, we will commonly refer to this condition as the cavity being in 'atomic resonance'. However, in any experimental setting, thermal drifts and vibration of the mirrors will occur, changing their separation. To sus-

tain atomic resonance in spite of environmental fluctuations, the mirrors are normally mounted on piezoelectric transducers. This allows their separation to be stabilised by electronic feedback, but requires a continuous measure of cavity length [45, 48]. As such, a frequency stable 'locking' laser is used to probe optical resonance, thereby establishing a range of spectral conditions on cavity design.

Firstly, the locking laser must be of sufficiently narrow linewidth to accurately resolve the resonance condition of the cavity. Then, in order to perform length stabilisation during atom-cavity interactions, the laser must be far detuned from atomic structure [56]. Otherwise, it may induce photon scattering events, prevent isolation of photons generated within the cavity and be routinely disturbed by vacuum Rabi splitting. This implies that it must address a neighbouring longitudinal resonance frequency of the cavity, separated by an integer number of free spectral ranges. Therefore $\Delta\omega_{fsr}$, and accordingly cavity length, must be tailored in consideration of the stable optical frequency references available. Alternatively, illumination and stabilisation of the cavity may be paused during the atomic interaction window using a 'sample and hold' technique. This requires the cavity to remain close to resonance when feedback is not being applied, placing demands on its mechanical stability and the maximum duration of atomic interaction. While this has proven an effective technique [120] requiring less experimental infrastructure, it constitutes a less elegant solution than off-resonant locking.

As the wavelength of the locking laser becomes further detuned from the design wavelength of the dielectric mirrors, it experiences a broader cavity linewidth. This decreases its responsivity with respect to cavity length fluctuations, establishing a limit on the bandwidth of the feedback signal it can generate. However, given that dielectric mirrors are only reflective over a well defined stop-band, the laser is ultimately limited to a modest detuning. At shorter cavity lengths, an increased free spectral range may cause consecutive longitudinal modes to observe large changes in cavity linewidth, greatly decreasing the number of viable locking frequencies. However, given the minimum

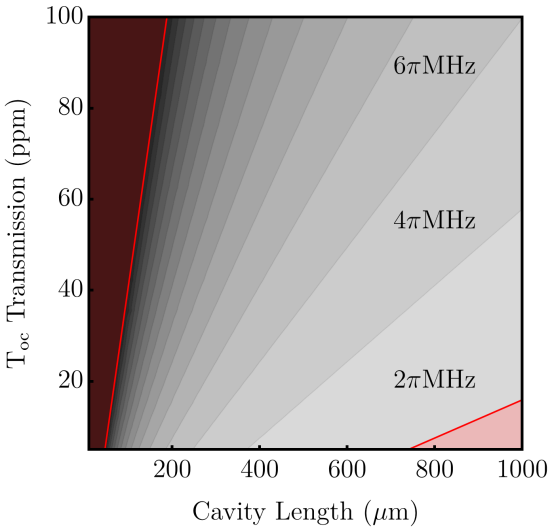

Figure 4.5: **Cavity Linewidth Limitation.** A contour map gives the variation of cavity linewidth, $\Delta\omega_c$, with cavity length and out-coupling mirror transmission. Each successive contour represents a linewidth change of 2πMHz, with three bands labelled directly. Standard mirror coating properties have been otherwise assumed, with $\mathcal{L}_{A+S} = 12$ppm and $T_{hr} = 2$ppm. Regions of cavity linewidth inappropriate for our experimental scheme are shaded in red. This includes the upper limit of $16 \times 2\pi$MHz, to assist the resolution of atomic structure in photon production. The minimum linewidth, 2πMHz, is derived from our method of cavity length stabilisation. This requires a frequency stable laser of narrower linewidth to prove cavity resonance, as discussed in Chapter 7.

cavity length required for optical access, this was not an issue experienced. A discussion of cavity locking is continued in Chapter 7, in the context of its experimental implementation.

4.5.2 Atomic Driving and Photon Production

We intend to use our optical cavities for single photon production, which as outlined in Chapter 2, requires an isolated interaction between the selected atomic transition and resonant mode of the cavity. If the cavity is resonant with multiple atomic transitions, or the atom with many longitudinal modes of the cavity, one may anticipate lower coherence[6] in the photons generated. To prevent this, we shall now outline a further

[6]This may be coherence with respect to self-interference in a g2 measurement, or with respect to a HOM measurement using an external source of single photons as reference.

set of limitations on the spectral properties of our cavity design.

Isolation of Atom-Cavity Interactions

In order for a cavity mode to be individually addressable, it must be well isolated in frequency from its neighbours. This implies that its free spectral range is significantly greater than its linewidth: equivalent to a statement of high finesse. Given our use of high reflectivity mirrors, this condition is arbitrarily guaranteed by Equation 3.3. Then, we require that the free spectral range is considerably larger than the atomic linewidth. For cavity lengths less than a millimetre, one can expect a free spectral range in the GHz regime. When compared to our standard atomic linewidth of $\Gamma \approx 6 \times 2\pi$MHz [121], this condition is again readily achieved.

Of more practical concern is to ensure that the linewidth of the cavity is sufficiently narrow as to address the intended atomic structure. However, given a common misconception of this statement, it is important to stress that the linewidth of cavity generated photons is not defined by the linewidth of the cavity itself. In similarity to standard atom-laser driving processes, photon frequency and bandwidth is given by energy conservation of the system and the timescale of atomic de-excitation. Indeed, it has been well demonstrated that the linewidth of cavity generated photons can be significantly narrower than that of the cavity [88]. In this way, the cavity simply acts as a spectral filter, limiting the frequency range over which atomic emission rates may be enhanced. Our maximum cavity linewidth derives instead from considerations specific to our atomic driving scheme.

Polarised Photon Production

A potential application of our developed atom-cavity system is to continue the development of a polarised single photon production scheme on the D2 line of ^{87}Rb [90]. This uses V-STIRAP to drive transitions between magnetic sublevels of the $^2\mathrm{S}_{1/2}\,|F = 1\rangle$ ground level, creating a stream of circularly polarised photons with alternating hand-

edness. As described in a recent thesis [56], an application of this work is to create a scalable source of single photons arriving simultaneously across many channels, for the development of linear optical quantum computing. This is achieved by polarisation dependent routing of the photon stream and applying a suitable time delay to each channel.

In order for the cavity to address the Zeeman split sublevels, the atom was subject to a magnetic field of sufficient strength as to create a splitting greater than cavity linewidth. However, for the cavity linewidth of $\Delta\omega_c = 3.75 \times 2\pi$MHz, splitting of $\Delta\omega_z = 20 \times 2\pi$MHz and corresponding field strength of $\approx$ 20G, the atomic level scheme itself was greatly disturbed. The Zeeman effect on the excited $^2P_{3/2}$ level in the D2 structure was placed into an intermediate regime, causing mixing of the hyperfine states. The full impact on single photon production is well described in a recent publication [122]. A simplified outcome is that the atom-cavity coupling corresponding to the two polarised transitions becomes asymmetric. In general, this causes one polarisation to be produced with a higher efficiency and purity than the other, which naturally proved detrimental to building a scalable source of single photon delivery.

The study [122] recommends a similar photon production scheme is driven on the D1 line of ^{87}Rb instead. The advantage comes in significantly greater hyperfine splitting of the $^2P_{1/2}$ level, preventing state mixing at the magnetic field strengths required for cavity interaction. However, this line comes with inherently weaker atom-cavity coupling, which for the aforementioned study, placed the existing cavity below the strong-coupling regime. In our following design work, we will not overly restrict cavity linewidth to optimise polarised photon production on the D2 line. However, the accessibility of a polarised photon production process will be an important consideration in deciding upon our mirror fabrication method, as well as being a notable application of the cavity constructed in Chapter 7.

Fundamental Mode Isolation

As our final spectral restriction, we must ensure that the atom interacts with only a fundamental Gaussian mode of the cavity. To facilitate remote entanglement, generated photons must have controlled spatial coherence. Photons that are transversely distinguishable are therefore not appropriate for establishing a coherent quantum interface. The fundamental Gaussian mode is the preferred spatial profile, given its compatibility with single-mode fibres. This allows the interconnection of distant cavity nodes without modal dispersion. However, photon transport within multi-mode fibres has been successfully applied to a number of demonstrative CQED experiments with fibre-tip cavities, as it ensures high efficiency geometric mode-matching [29].

To ensure isolation of the fundamental mode of interest, we must verify that the free spectral range of the cavity is not close to an integer multiple of its higher order mode spacing. Given that these both depend on cavity geometry, there exists a large set of cavity parameters that may cause co-resonance of the atom on two modes. However, as the order of the transverse mode increases, so too does its mode volume, decreasing atomic coupling strength. Further, many higher order modes have zero intensity along the cavity axis, the intended site of atomic localisation. In most practical scenarios, it therefore proves sufficient to consider only the position of a few lower order modes. As was asserted in Section 3.3.4, the geometrical region around confocality must certainly be avoided, as even order transverse modes coincide exactly with fundamental modes.

4.6 Design Optimisation

In determining the experimental requirements of our cavity, we are now able to perform the optimisation of its design. First, we derive general metrics of performance for its role in single photon production. This shall be done for two and three level atomic driving schemes, allowing for a more quantitative comparison of their standard efficiencies. We consider the extent to which previous cavity systems have met our ongoing experimental

challenges and motivate the route towards improvement. This clarifies our motivation for migrating to the use of micro-mirrors, with laser ablation chosen as the platform for their fabrication. Finally, we consider a somewhat alternative approach to high efficiency single photon production, by engineering birefringence into our cavity.

4.6.1 Elementary Performance Metrics

To allow for the numerical optimisation of our cavity design, a quantitative metric for its performance must be developed. To begin, we return to the simple model of the two-level atom derived in Section 2.4.1. We assume that a rapid excitation process is applied to the atom, such as a strong π pulse, placing it suddenly into an excited state. A single photon is then emitted either into the cavity mode, or free space, with the former establishing our photon source.

The two level atom consists of an excited $|x\rangle$ and ground $|g\rangle$ atomic state, whose transition frequency is resonant with an optical cavity. The wave function of this system is given by:

$$\Psi(t) = c_x(t)\,|x,0\rangle + c_g(t)\,|g,1\rangle\,, \tag{4.4}$$

where $c_i(t)$ is the probability amplitude of the system in state i at time t. The evolution of this system is solved by a time-dependent Schrödinger equation, assuming an initially excited atom. Then, the efficiency of photon generation, η_{gen}, can be determined by the probability of the atom being in its ground state, multiplied by the single photon decay rate:

$$\eta_{gen} = 2\kappa \int_0^\infty |c_g(t)|^2 dt' = \frac{\kappa g^2}{(\kappa+\gamma)(\kappa\gamma+g^2)} = \frac{\kappa}{\kappa+\gamma}\frac{2C}{2C+1}, \tag{4.5}$$

where we have reintroduced system cooperativity, C. A common simplification is to consider the *fast cavity regime*, where $\kappa \gg \gamma$. Using the form of Purcell enhancement derived in Equation 2.19, photon generation efficiency is then given by:[7]

$$\beta_{gen} = \frac{W_{cav}}{W_{cav} + W_{free}} = \frac{F_p}{F_p + 1} = \frac{2C}{2C + 1}. \tag{4.6}$$

Outside the fast cavity regime, β_{gen} neglects a significant portion of the excitation lost to atomic depolarisation. This occurs if the atom remains in its excited state for a prolonged period of time, due to low atom-cavity coupling, g. Or, atomic depolarisation may also be substantial for very high g, as re-excitation of the atom by the photon will occur faster than it decays via one of the mirrors. It therefore follows that the optimal value of η_{gen}, for a given g and γ, is reached when $\kappa = g$; equivalent to the critical damping of Rabi oscillations.

To produce photons useful for quantum information processing, photon decay must be in transmission from the out-coupling mirror. Thus of the total photonic decay rate, κ, we must consider its fraction into this coherent channel, κ_{out}, with respect to incoherent channels, κ_{loss}; such that $\kappa = \kappa_{out} + \kappa_{loss}$. This allows us to define a total 'output efficiency' of the cavity:

$$\eta_{out} = \beta_{ext}\eta_{gen} = \frac{\kappa_{out}}{\kappa}\eta_{gen} = \frac{g^2}{g^2 + 2\kappa\gamma}\frac{\kappa_{out}}{\kappa + \gamma}. \tag{4.7}$$

In the above expression, we have separated output efficiency into two terms. The first considers atomic emission into the cavity, expressing cooperativity in terms of its constituent rates. The second encodes photon extraction efficiency, where out-coupling

[7] Note the subtle assertion that the emission rate of the atom into free space, W_{free}, is equal to the atomic emission rate outside of the cavity. As discussed in Section 2.3.1, this is because for the cavities considered in this thesis, the solid angle subtended by the mirrors is small. As such, the atom is not shielded from its environment and continues to spontaneously emit into it at a normal rate.

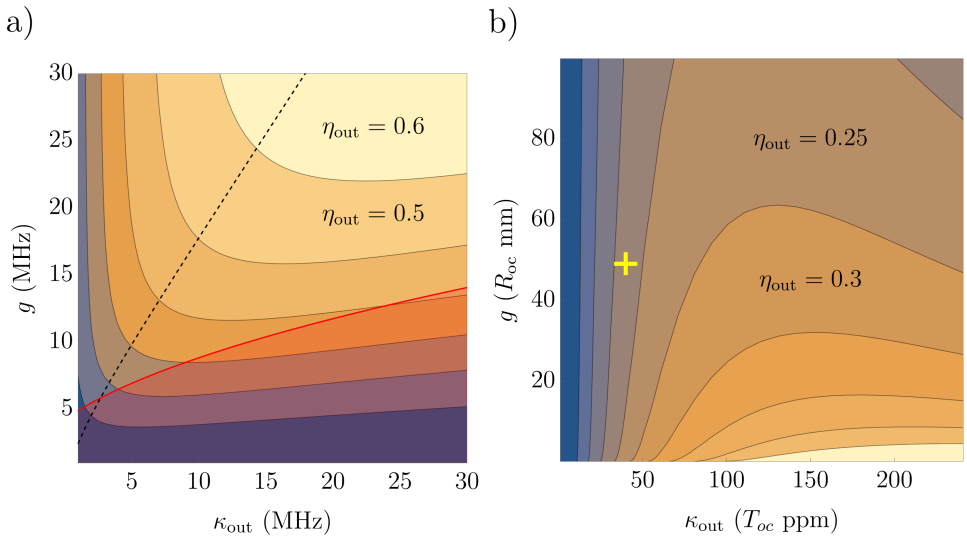

Figure 4.6: **Photon Output Optimisation. a)** A contour map of η_{out} is shown, with respect to varying g and κ_{out}. There are two regions of interest, where each parameter takes dominance over the process efficiency, separated by a dashed line. Also shown in red is the region of low cooperativity, $C < 1$. We assume here that $\gamma = 3$MHz and $\kappa_{loss} = 3$MHz. **b)** Shown is a further contour map of η_{out}, but with g and κ_{out} being varied by their respective physical parameters R_{oc} and T_2. Remaining cavity parameters are taken to be constant, given by the 2014 cavity described in Table 4.1. The exact position of this cavity is shown as a yellow cross, highlighting that it exists within a regime where increased g is of little improvement to process efficiency alone.

is compared to unfavourable decay channels, γ and κ_{loss}. To examine which of these terms is dominant in the optimisation of photon output efficiency, we have plotted the relationship of g and κ_{out} to η_{out} in Figure 4.6.

As is demonstrated by this figure, there are two clear regions of optimisation, separated by $\nabla_g \eta_{out} = \nabla_{\kappa_{out}} \eta_{out}$. In the lower region[8], the output efficiency of the system is limited by atom-cavity coupling, requiring a further reduction of cavity mode volume. Above this line, it is limited by photon extraction efficiency, requiring greater transmission of the out-coupling mirror. As g and κ_{out} both depend on L_{cav}, it has been fixed for this plot. However, the length of the cavity is clearly a key parameter in its optimisation, as discussed later. To highlight that these parameters depend non-

[8]Given by $4\gamma^2\kappa\kappa_{out} + 4\gamma\kappa^2\kappa_{out} > g\big((\gamma + \kappa_{loss})(g^2 + 2\gamma\kappa_{loss}) - 2\gamma\kappa_{out}^2\big)$.

linearly on their associated physical parameters, Figure 4.6 varies g and κ_{out} via R_{oc} and T_{oc} instead.

4.6.2 Comparing Photon Production Schemes

From the perspective of single photon production, the two-level system has some clear drawbacks. Primarily, the presence of atomic re-excitation is highly detrimental to photon output efficiency, as it increases the timescale over which atomic depolarisation may take place. Indeed, even for perfect mirrors, $\kappa = \kappa_{out}$, unit photon output efficiencies are not readily achieved. This is rather disturbing from the perspective of cavity design, as perfect mirrors are impossible to fabricate.

A key aspect of photon production by V-STIRAP, and a motivation for its adoption, is the mitigation of γ by only transiently populating an excited atomic state. This allows for near unit photon generation efficiencies, in cavities where $\kappa < \gamma$. This greatly simplifies cavity design, where κ is not tailored by suppression of Rabi oscillations, but by the control of cavity linewidth, cooperativity and photon extraction efficiency. In this way $\beta_{out} = \beta_{gen}\beta_{ext}$ proves an accurate reflection of cavity performance, supported by a number of studies [19].

Whilst single photon production is certainly one demonstration of a versatile atom-cavity interface, V-STIRAP has further advantages over the two-level scheme. These aspects have been discussed in Chapter 2, but include an inherent reversibility of the photon production process and the ability to encode quantum information within their spatio-temporal wave-packets. However, given the adiabatic following required of the Raman process, it may be limited in the production of a high bandwidth photon source.

4.6.3 Existing Cavity Analysis

We have now developed a comprehensive set of tools to assist with the quantitative analysis of cavity design and its optimisation. To begin, we will examine two genera-

tions of optical cavities produced within our research group, formed from super-polished mirrors. We will retrospectively evaluate their suitability for single photon production and motivate a direction in which improvements can be made.

Our first cavity was produced by J. Dilley and P. Nisbet-Jones in 2010 [28, 123], with geometric and spectral properties given as Table 4.1. This used near-planar mirrors of $R_{oc} = 50$mm, standard for super-polished substrates at the time. To reduce mode-volume and enhance atom-cavity coupling, the length of the cavity was relatively short at $L_{cav} = 74\mu$m. Some contamination of the mirrors occurred during their installation into vacuum, increasing incoherent mirror losses from 2ppm to 17ppm. When combined with its short length, the linewidth of the cavity became $\Delta\omega_c = 24 \times 2\pi$MHz. This was considered too broad for polarised single photon production, given the magnetic field strength required to resolve Zeeman split sub-levels. Then, since the optical access was considerably low at $\theta_{oa} = 4^\circ$, it was unsuitable for intra-cavity dipole trapping.

A second cavity was produced by T. Barrett in 2014 using an identical set of mirrors [56]. In order to be suitable for polarised photon production, cavity length was increased to 339μm and further measures were taken to prevent contamination of the mirrors. This gave a linewidth of $\Delta\omega_c = 3.75 \times 2\pi$MHz, now well within our defined limit. The lower losses of the mirrors also enhanced β_{ext}, thereby surpassing the photon output efficiency achieved in 2010. However, with its optical access remaining low at $\theta_{oa} = 18.5^\circ$, it would remain unsuitable for intra-cavity dipole trapping.

To address the issue of optical access, further coning was performed on the mirrors, creating a final facet diameter of 300μm. A reproduction of the 2014 cavity would then, in principle, fulfil all of our stated requirements. However, to evaluate if some improvement could be made to this design, we now perform its numerical optimisation. A grid search algorithm sampled coordinates of (L_{cav}, T_{oc}), optimising β_{out}. The parameters $(R_{oc}, \sigma_m, \sigma_s, T_{hr}, \mathcal{L}_T)$ were held fixed, using the mirror properties of 2014. All potential

Property	Symbol	Year		Optimisation		Unit
		2010	2014	SPol	Micro	
Cavity Length	L_{cav}	74	339	310	400	μm
Radius of Curvature	R_{oc}	50	50	50	0.25	mm
Mirror Diameter	σ_m	1	1	0.3	0.06	mm
Transmission (OC)	T_{oc}	40	40	71	240	ppm
Facet Diameter	σ_s	1	1	0.3	0.13	mm
Transmission (HR)	T_{hr}	0.1	0.1	0.1	2	ppm
Mirror Losses	$\mathcal{L}_{A+S}$	17	6.5	6.5	12	ppm
Cooperativity	C	5.2	3.4	2.2	20.6	-
Cavity Linewidth	$\Delta\omega_c$	24	3.75	6.5	16	2π MHz
Optical Access	θ_{oa}	4	18.5	45	72.5	$^\circ$
Coupling Strength	g	19.3	6.2	6.6	30.6	2π MHz
Photonic Decay	κ	11.94	1.9	3.24	7.9	2π MHz
Finesse	$\mathcal{F}$	85,000	118,000	75,000	23,000	-
Generation Efficiency	η_{gen}	0.73	0.33	0.42	0.70	-
Output Efficiency	η_{out}	0.39	0.25	0.36	0.66	-
Generation Efficiency	β_{gen}	0.91	0.87	0.82	0.98	-
Extraction Efficiency	β_{ext}	0.54	0.75	0.84	0.94	-
Output Efficiency	β_{out}	**0.49**	**0.66**	**0.69**	**0.92**	-

Table 4.1: **Optical Cavity Performance Analysis.** The properties and expected performance for a range of optical cavities is described. This includes two systems used for quantum information processing within our research group, fabricated in 2010 [28] and 2014 [56]. Then, by numerical design optimisation, we indicate the ideal cavities that can be produced using super-polished mirrors and micro-mirrors. The properties of each system is specified in sections, consisting of their free design parameters, fixed parameters, fulfilment of experimental requirements (coloured in red when unmet and orange for the limit of fulfilment), standard comparative metrics, efficiency for standard Purcell enhanced photon capture and for V-STIRAP.

cavity designs were ensured to meet our spectral and geometric requirements, specifically: $C > 1$, $1 < \Delta\omega_c/(2\pi\text{MHz}) < 16$ and $\theta_{oa} > 45^\circ$. The ideal cavity design indicated is given in Table 4.1, making a modest improvement to photon output efficiency using V-STIRAP.

Interestingly, photon generation efficiency is now lower than in previous designs, with the overall improvement derived purely from their increased extraction. From the form

of β_{gen} given in Equation 4.6, we can see that near unit efficiency is achieved for $C \gg 1$. Accordingly, increasing β_{out} by only a reduction in mode volume quickly saturates any efficiency gain. However, it does provide freedom to then increase mirror transmission, keeping system cooperativity largely constant. That is, any gain to g is optimally balanced out by increasing κ_{out}. In this optimisation, we have reduced cavity length by 29μm, to its minimum value for 45° optical access. Then, we have increased mirror transmission by 31ppm. While this gives only a modest increase to photon output efficiency, it introduces a key design principle of micro-cavities and sets a benchmark of performance for super-polished mirrors.

4.6.4 Migration to Micro-Mirrors

In the optimisation of super-polished mirror cavities, we were unable to further increase atom-cavity coupling due to our requirement for optical access. As a result, there are two practical solutions to achieve greater photon output efficiency. The first is to reduce mirror facet diameter, σ_s, to allow for greater freedom in cavity length selection. Or, instead of reducing cavity mode volume via cavity length, we can instead increase mirror curvature. From a technical perspective, each of these options suggests a migration towards the use of micro-mirrors in cavity fabrication. In this section, we shall quantitatively analyse the benefits of micro-cavities and decide on their ideal method of production.

We now re-perform cavity optimisation, expanding our tailored degrees of freedom to the full range outlined in Section 4.2. Our sampled coordinates of mirror geometry, (R_{oc}, σ_m) is now over the range of micro-mirrors defined in Section 4.3.4. For each of these coordinates, the ideal values of L_{cav} and T_{oc} are chosen by a grid search optimisation of β_{out}. Following our discussion of laser ablation in Section 4.3.3, we consider mirror facet diameter to be that of a single-mode optical fibre, $\sigma_s = 125\mu$m. We then assume optimistic, but realistic, coating properties of $T_{hr} = 2$ppm and $\mathcal{L}_T = 12$ppm. All cavities are verified to fulfil our requirements of strong-coupling, optical access and

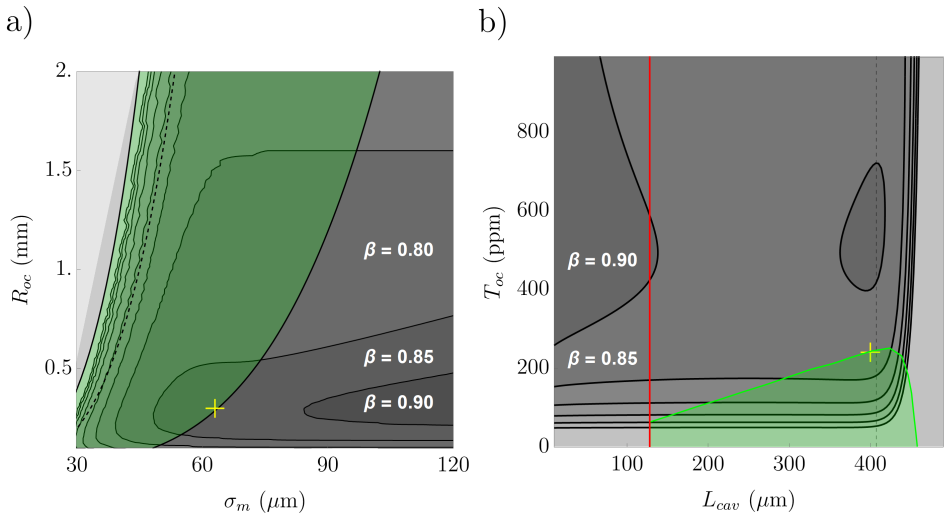

Figure 4.7: **Cavity Optimisation Process. a)** The optimisation of β_{out} was performed, where each coordinate of (σ_m, R_{oc}) represents a grid search optimisation of (L_{cav}, T_{oc}). The green region denotes the standard geometrical regime of ablated micro-mirrors, following our analysis of Section 4.3.4. The ideal cavity coordinates are indicated by a yellow cross and are at the limit of the ablated region. Improved cavity performance is possible for mirrors of greater diameter. **b)** For the ideal values of R_{oc} and σ_m, the optimisation of T_{oc} and L_{cav} is shown. The region where our experimental requirements are met is shown in green. Otherwise, left of the red line represents insufficient optical access and right of the line denotes where cavity linewidth is too broad. At the greatest cavity lengths the onset of clipping loss can be seen, rapidly reducing photon output efficiency.

spectral suitability as outlined in Sections 4.4 and 4.5. The properties of the ideal cavity indicated are given in Table 4.1.

The use of micro-mirrors has seen a substantial improvement to both β_{gen} and β_{ext}. Their high curvature has greatly decreased cavity mode volume, increasing cooperativity to $C > 20$. This has essentially saturated photon generation efficiency at $\beta_{gen} = 0.98$. As introduced in our optimisation of super-polished cavities, this has allowed T_{oc} to be increased by a substantial 170ppm, while remaining in the strong-coupling regime. Thus despite the greater incoherent losses expected of micro-mirrors, its negative impact on extraction efficiency has been rendered negligible by greater κ_{out}.

The geometrical mirror parameters of our optimal cavity are $R_{oc} = 250\mu$m and $\sigma_m = 60\mu$m. From our comparison of glass shaping methods given as Section 4.3.4, this cements laser ablation as our chosen method of their production. To now guide the design of this experimental platform, highlighting which mirror properties should assume priority, we examine this optimisation further. Accordingly, Figure 4.7 shows regions of β_{out} varied around our selected cavity parameters.

In this figure, it is clear that optical access and high cooperativity are no longer a limiting factor to cavity design. Instead, we are limited by the maximum linewidth of our cavity, at $16 \times 2\pi$MHz. Ideally, we would be able to narrow cavity linewidth by a further increase in cavity length. However, we are currently limited in this regard by the onset of clipping losses. As is discussed in Section 4.4, the most practical solution to this is to simply increase mirror surface diameter. In this effort, our experimental platform will use a recently developed method of laser ablation, termed 'laser-dot milling' [47]; demonstrated to enhance feature diameter. It then remains important to achieve good optical quality of the final mirror surfaces, as greater incoherent losses would frustrate our efforts to reduce cavity linewidth.

Finally, the higher standard cooperativity of cavities formed from micro-mirrors increases our experimental versatility. We have discussed that for our favoured scheme of polarised photon production, a migration to the D1 line would promote higher photon production efficiency and purity. However, to overcome the inherently weaker atom-cavity coupling when compared to the D2 line, a more confined mode volume is required. Indeed, this was the reason the optimised super-polished cavity was never fabricated. Although it met with our general requirements, the far greater freedom in cavity applications that comes with micro-mirrors is preferable.

Outside the context of single photon production, many quantum information processing schemes require several coherent operations to be performed on the emitter. As coop-

erativity is defined as a ratio of coherent to incoherent dynamic rates inside the cavity, it is also the basic figure of merit for the number of quantum operations that can be performed. Within our field, cavities formed from micro-mirrors have realised in excess of $C = 100$, largely impossible with open-access super-polished substrate cavities [124].

4.6.5 Birefringent Orthogonal State Coupling

Here, we consider a potential advantage of engineering birefringence into our cavity mirrors. This discussion is based on our recent work [125], for which I was a co-author. An expanded discussion of the technique and full list of contributors can be found therein. We start by taking a simplified form of photon generation efficiency in the two level atom. For the case of high cooperativity, Equation 4.5 reduces to:

$$\eta_{gen,C\gg1} = \frac{\kappa}{\kappa + \gamma}. \tag{4.8}$$

This sets our basic measure of process efficiency for the non-birefringent case. A birefringent cavity supports two orthogonal polarisation eigenmodes, separated in frequency by $\Delta\omega_\varepsilon$ [38]. Some degree of birefringence is inherent to mirrors comprised of dielectric stacks [126], but is also engineered into a cavity by the use of geometrically elliptical mirrors [39]. When light circulating the cavity is misaligned with respect to a polarisation eigenmode, its polarisation state rotates at a rate $\Delta\omega_\varepsilon$. This equally applies to atomic emission into the cavity: where an atom emitting into a non-eigenmode of the cavity will generate a photon in an evolving polarisation state.

To consider the potential utility of this process, we define a birefringent cavity with linearly polarised eigenmodes, $|X\rangle$ and $|Y\rangle$. This is coupled to an atom which may emit circularly polarised light, $|+\rangle = 1/\sqrt{2}(|X\rangle + i\,|Y\rangle)$ or $|-\rangle = 1/\sqrt{2}(|X\rangle - i\,|Y\rangle)$, as quantised in the axis of the cavity. The atomic level scheme, described in both the atomic and cavity basis, is given as Figure 4.8. We consider each linear polarisation

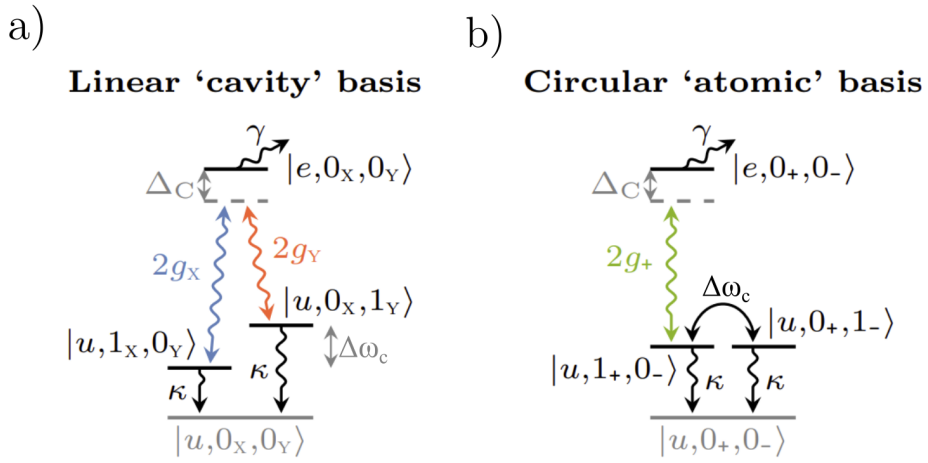

Figure 4.8: **Birefringent Cavity and Two Level Atom. a)** In the basis of the birefringent cavity, the atom emits into both horizontally and vertically polarised eigenmodes. The energetic splitting of these states is $\Delta\omega_c$. Each mode is subject to decay, via coherent emission of the photon from the cavity. **b)** In the atomic basis, the atom couples only to one circularly polarised mode, with the photon evolving to the uncoupled state at a rate $\Delta\omega_\varepsilon$. This reduces re-excitation of the atomic state, $|e,0_+,0_-\rangle$, decreasing population lost via γ. Figure adapted from [125].

mode to have equal detuning $\pm(1/2)\Delta\omega_\varepsilon$ from atomic resonance.

When an atom emits into $|+\rangle$, its linear components experience a progressive dephasing, induced by the birefringence of the cavity. This will evolve the photon towards a $|-\rangle$ state, thus decoupling it from the atom. As a greater photon population is de-coupled, re-excitation is prevented and atomic de-polarisation is suppressed. Essentially, birefringence has allowed us to shield the photon from the atom, enhancing its extraction efficiency.

It can be shown, via the quantum master equation approach taken in the paper, that an optimal choice of birefringence is $\Delta\omega_\varepsilon = \sqrt{2}g$. This doubles the time in which the atom is in its ground state, providing an equal enhancement to the effect of κ on photon generation efficiency. Thus, we re-express Equation 4.8 as:

$$\eta_{gen,C\gg1,\Delta\omega_\varepsilon=\sqrt{2}g} = \frac{2\kappa}{2\kappa+\gamma}. \tag{4.9}$$

In the case of the 2010 cavity, this increases η_{gen} from 0.73 to 0.81 and for the 2014 cavity from 0.33 to 0.49. For $C \gg 1$, the maximum enhancement that can be engineered by this method is 17.2 percentage points, achieved for $\kappa = \gamma/\sqrt{2}$. However, we are ultimately limited in this method by the continuous evolution of the polarisation state of the photon, which cannot simply be halted once decoupled from the atom. To reach unit photon generation efficiencies, one requires access to a larger number of states which are decoupled from atomic de-excitation.[9]

Overall, this highlights birefringence as an interesting degree of freedom in cavity design, used in the control of polarisation dependent photon production schemes. In the following work on mirror fabrication, its control is considered in detail, by engineering cavity mirrors with an elliptical geometry. However, for schemes outside the production of a single photon, the evolving polarisation state of the photon must be treated carefully. In particular, it may generate a limitation on the unification of many single photon sources [38], or reduce the fidelity of remote atom entanglement within a quantum network [127].

4.7 Design Review

The central objective of this chapter was to consider the implications placed on cavity design by a requirement to facilitate strong atom-cavity coupling to a dipole trapped single atom. We have outlined the geometrical limitations this creates, in ensuring sufficient open-access to the cavity mode for injection of the trapping light. However, since the properties of a cavity as a coherent quantum interface have a highly interrelated dependence on its small set of physical design variables, we have observed that

[9] The paper describes this scenario using an N-level atom.

achieving this open access must be performed with care.

In this effort, we evaluated previous atom-cavity systems and considered their limitations. We determined that an increased mirror curvature allows for longer cavity lengths and increased mirror transmission, without leaving the strong-coupling regime. In simultaneously promoting photon extraction efficiency, system cooperativity and optical access, the route towards creating our desired cavity was with the use of micron-scale cavity mirrors. However, the diameter of these mirrors proves an important aspect of preventing clipping losses, allowing cavity length to be increased in favour of a reduced cavity linewidth or improved optical access.

In the opening of this chapter, we suggested that a cavity highly optimised for one experimental architecture may have limited versatility. By this reasoning, we favoured more elementary metrics of optical cavity performance. Accordingly, in the work of the following chapter we seek to more broadly expand the range of accessible mirror geometries, rather than pursue the construction of only a single cavity schematic. However, by considering some more refined examples of single photon production, we have deduced where the emphasis of this work should be placed. In highlighting a previous study of producing polarised single photons on the rubidium D2 line, we noted that an inherently larger cavity cooperativity provides freedom to move to a more weakly coupled atomic transition. In examining a route towards increased photon production efficiency in a two level atom, we have observed a potentially significant advantage of controlling cavity birefringence.

Overall, we intend to produce mirrors of tight curvature, with large diameters and controlled ellipticity. In consulting the standard geometrical range of FIB milling and laser ablation given in Section 4.3.4, this can be achieved directly using ablation. As discussed therein, this comes with the immediate advantage of good optical surface quality, simplifying the creation of a high finesse resonator. However, unlike most

approaches, we will not be seeking to create fibre-tip cavity mirrors. In our process of cavity design, we have treated our micro-mirrors as standard optical elements, without needing to consider the overlap of cavity mode geometry with that guided by the fibre. Doing so would likely dominate our cavity design optimisation process, or require the development of complex optical fibres with inherent geometric mode-matching. We therefore seek to create a mirror which has the same facet geometry as a fibre-tip, but the experimental practicality of a bulk mirror substrate. This allows propagation of the cavity mode into free space, without further restrictions.

Chapter 5

Cavity Mirror Fabrication

This chapter describes a novel apparatus for laser ablation, used in the production of micron-scale cavity mirrors. We open with a technical description of the developed platform, including measures taken to control the ablative process and reconstruct fabricated surface features. Then, it will be applied to the manufacture of pyramidal micro-mirrors, evaluating their role as a potential alternative to fibre-tip cavities.

5.1 Introduction to Laser Ablation

Laser ablation is the removal of material from a substance by evaporation or sublimation, thermally induced by intense laser irradiation. It is routinely used in material processing to produce detailed micro-structures on materials that are otherwise difficult to machine. Sub-microsecond pulsed lasers are a common choice for this application, given their ability to impart high energies in a controlled manner [128]. However, they have proven unsuitable for the processing of optical glass, given the relatively low surface quality of their generated features [129, 130].

In order to create spherical depressions suitable for cavity mirrors, continuous wave (CW) lasers are used instead, with exposure times of greater than 1ms [36]. This creates a molten layer underneath the evaporated material, which due to surface tension, solidifies with high optical quality. The interplay of evaporated and molten glass can be controlled by the characteristics of the laser exposure, allowing one to shape the feature created. However, the control of fast and complex thermal processes is notoriously difficult, imposing many limitations on the geometry and reproducibility of

mirror fabrication [115].

The majority of research into laser ablation has been towards establishing greater control over this thermal process, allowing the fabrication of a wider range of surface geometries with fewer errors. This is highlighted by the development of several methods of expanding feature diameter and controlling their ellipticity [40, 47, 131]. However, an analytical model has not yet been developed which can directly connect laser pulse parameters to produced geometries [132]. Therefore, a detailed empirical characterisation of each ablation system and substrate design must be performed regularly, rendering the creation of arbitrary mirror profiles highly time consuming. This requires that increasingly sophisticated methods of mirror shaping be matched by advances in rapid surface reconstruction.

5.1.1 Experimental Objectives

Our discussion on cavity design proposed the advantages of micro-mirrors for building a quantum interface, via an increased photon output efficiency and allowing for greater experimental flexibility. Therefore in our construction of an ablation system, we are primarily looking to access the standard geometric regime of ablated mirrors, presented in Section 4.3.4. Then, we wish to concentrate on two improvements, highlighted by our examination of cavity optimisation. The first is to create larger mirror surface diameters, preventing the onset of clipping losses at extended cavity lengths. This allows for greater optical access and a reduced linewidth. Secondly, we wish to control mirror ellipticity, allowing a further degree of freedom in the control of polarisation dependent driving schemes.

This enhanced geometric control will be achieved by the combination of two established methods, micro-patterned ablation [47] and substrate rotation [40], each reliant on precise substrate translation. Given some technical issues presented by fibre-tip cavity mirrors, we will additionally be looking to perform ablation on a variety of prototype

substrates, requiring an adaptable substrate mount. Then, in order to overcome the laborious nature of connecting feature geometry to laser pulse parameters, we require a rapid and non-invasive method of surface reconstruction.

In fulfilling these requirements, we also open the novel possibility of building arrays of micro-mirrors placed on the same substrate. We investigate this for two applications. The first is to increase the scalability of single atom-cavity systems, allowing multiple quantum interfaces to be operated using the same experimental architecture.[1] This application is further discussed in Section 6.5, in the context of an initial demonstration of a four-cavity system. Secondly, an array of several hundred features may be used for enhanced chemical sensing. While outside the scope of this thesis, several studies support optical cavities as a useful tool in the rapid analysis of chemical samples [133, 134].

5.1.2 Overview of Platform

Our experimental approach to mirror shaping was to overlap the beamlines necessary for laser ablation and surface profilometry. The first required an infrared beamline, to apply precise and stable laser exposures to the target substrate. The second comprised a white light interferometer, to monitor evolving substrate geometry. In order to rapidly transition between these operations, their required optical elements were mounted on motorised flip mounts. An electronically re-configurable approach allowed glass shaping to proceed with greater automation, a design intention which was reflected in many aspects of the developed platform. An illustration of the combined system is given as Figure 5.1, showing both configurations. In the remainder of this section we shall highlight the operating principles of this system, with detailed technical descriptions given in the remainder of the chapter.

[1] For the cavity systems considered in Chapter 7, this would be a common vacuum chamber and system for length stabilisation.

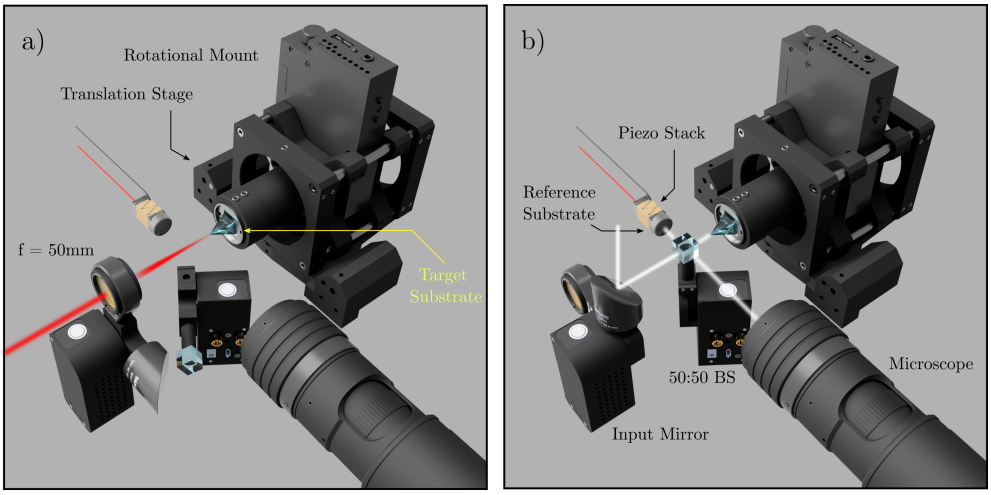

Figure 5.1: **Ablation and Profilometry Integration.** The developed system overlaps the optical elements necessary for glass shaping and surface reconstruction, via the use of motorised flip mounts. **a)** For laser ablation, a collimated infrared beam is focused onto the target substrate by an f = 50mm lens. The target substrate is mounted on a rotational and translational stage. **b)** A Michelson interferometer is used for surface reconstruction with white light illumination. This requires introduction of a 50:50 non-polarising beam splitter and planar reference substrate. The optical path lengths to the reference and target substrates are tuned by piezoelectric stacks, with their joint interference pattern measured by microscope.

During the ablation process, the target substrate is exposed to an illumination power of 1-4 watts for a duration of 5-150 milliseconds. These pulses are shaped by the rapid modulation of laser excitation, which was observed to give an improved performance over using a fast shutter. Electronic feedback is applied to reduce fluctuations in intensity, improving the geometrical consistency of generated features. To then average out any residual error, each pulse sequence comprised a large set of repeated exposures.

Glass shaping and profilometry required continual substrate re-positioning, provided by motorised translation stages. Linear stepper motor translations were used to create large arrays of features on one substrate. Piezoelectric actuators were used to apply patterned exposures and perform fine tailoring of individual feature geometry. Substrate rotations were applied to smooth feature ellipticity, or set the orientation of an elliptical feature with respect to its base substrate.

Feature geometry was determined by white light interferometry, via phase and coherence scanning methods. This provided a measure of feature diameter, curvature and ellipticity; allowing the empirical connection of a pulse sequence and its associated mirror geometry. To evaluate surface roughness and predict scattering losses, atomic force microscopy was performed external to the ablation apparatus.

All elements were combined by a unified software control. This allowed the fabrication of an array of features and their characterisation to be performed with minimal user input. A full schematic of the experimental system is given as Figure 5.2, which shall be referenced throughout the remaining chapter.

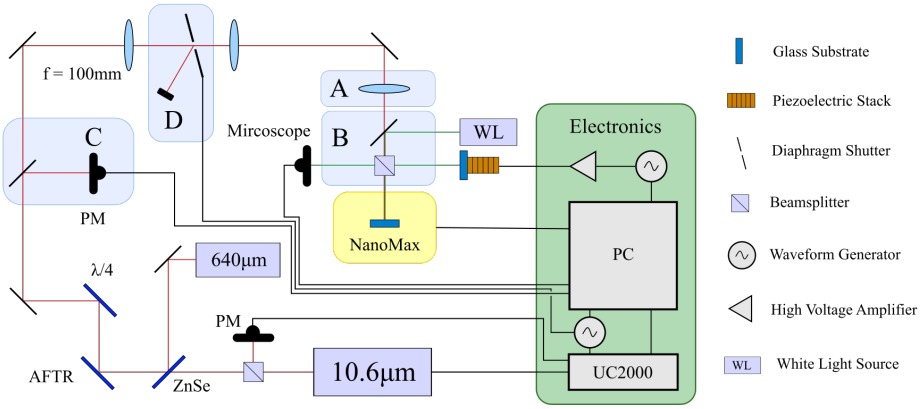

Figure 5.2: **Schematic of Laser Ablation and Surface Reconstruction System.** The combination of 10.6μm, 640nm and white light beam paths is shown. The general wiring diagram is outlined, highlighting which apparatus is directly synchronised. The electronically removable sections of the apparatus are lettered, which allow the system to rapidly switch between A: ablation and B: interferometry. Ablation pulse shaping is assisted by C: power stabilisation using a power meter (PM) and D: fast control of substrate exposure by a shutter. A key is given to the right of the schematic. Speciality optics include an absorbing thin film reflector (AFTR), quarter wave retarder ($\lambda/4$), and zinc selenide (ZnSe) window used as a dichroic mirror.

5.2 Ablation Laser System & Pulse Generation

The nonlinear nature of ablation renders feature geometry highly sensitive to laser power and exposure duration. As such, these must be well controlled in order to achieve feature regularity and reproducibility. Since peak absorption for glass lies in the 10μm wavelength regime, relatively inexpensive CO_2 lasers are used. However, these lasers are generally supplied for industrial cutting applications, so high stability is not present off-the-shelf. Instead, bespoke systems for regulating power fluctuations must be constructed in the laboratory. Accurate pulse shaping additionally proves more challenging at this wavelength, due to the lesser availability of fast electro-optical devices. As we shall discuss, some interplay between improving stability and accurate pulse-shaping was required in our approach.

5.2.1 Laser System

The laser selected was a Synrad 48-1w with a maximum optical power of 10W at 10.6μm. This is normally a pulsed laser, whose power is determined by the duty cycle of a 5kHz driving signal. However, in increasing this repetition rate to 20kHz, the laser behaves quasi-continuously over the timescale of ablation. The collimated output beam diameter of the laser is 3.5mm, focused onto the target ablation substrate using an f = 50mm air spaced doublet lens.[2] The temperature of the laser is stabilised using a closed-loop centrifuge water cooler,[3] which maintains an operational temperature of $\pm$ 0.1°C up to maximum output laser power. Whilst air-cooling is standard for gas lasers, water cooling gives greater thermal stability.[4]

Back-reflections from beam-shaping optical elements into the lasing medium are likely to cause significant power instability. Therefore, an optical isolator was constructed on

[2] Thorlabs AC254-050-F

[3] PolyScience LM6

[4] This also conveniently allows fan vibration to be located away from the optical table surface, preventing oscillation of attached optical elements.

the output of the laser. This primarily consisted of an absorptive thin film reflector (AFTR), which reflects s-polarised light and absorbs p-polarised light, followed by a quarter wave retarder (QWR).[5] Light output from the laser is s-polarised, so is reflected by the AFTR and rotated into circular polarisation by the QWR. Previous studies have reported lower feature ellipticity using circularly polarised light [39], therefore we do not modify it further. Light back-reflected onto the QWR is rotated into p-polarisation and absorbed by the AFTR, thus protecting the laser. Since the operation of the AFTR requires it to be mounted at 45° to the direction of propagation, the optical isolator is additionally used as a fixed periscope to match the height of the target substrate.

Even with a temperature controlled environment, the intensity stability of the laser is approximately ±5%, which is unsuitable for precise ablation. Therefore, systems of fast and slow feedback were developed. High frequency fluctuations are measured using a 90:10 beam splitter on the laser output and monitoring of the weaker channel using a thermopile detector.[6] An error signal is fed directly into the laser controller RF sequencer to modify the duty cycle of the laser. However, it can be seen that discontinuous power jumps, corresponding to modal competition between the 10.53μm and 10.67μm laser lines in CO_2, are unaffected. This behaviour was readily seen in fabricated cavity arrays by the presence of two distinct feature geometries.

5.2.2 Active Power Stabilisation

In order to reduce mode-hops and slow drifts in laser power, the intensity is further monitored using a thermal power meter[7] between ablation pulses, with an integration period of 1s. The required power corrections are of sufficiently low bandwidth that they can be made by software adjustments to the duty cycle. This applies a proportional-integral-derivative (PID) feedback signal to the laser controller, allowing the user to 'lock' the laser power to a desired value. This does not require one to calibrate the rela-

[5] Ophir 631347-117 & 630706-117
[6] Synrad Closed Loop Stabilization Kit
[7] Thorlabs S425C

tionship between duty cycle and output power, which measurably shifts on a daily basis.

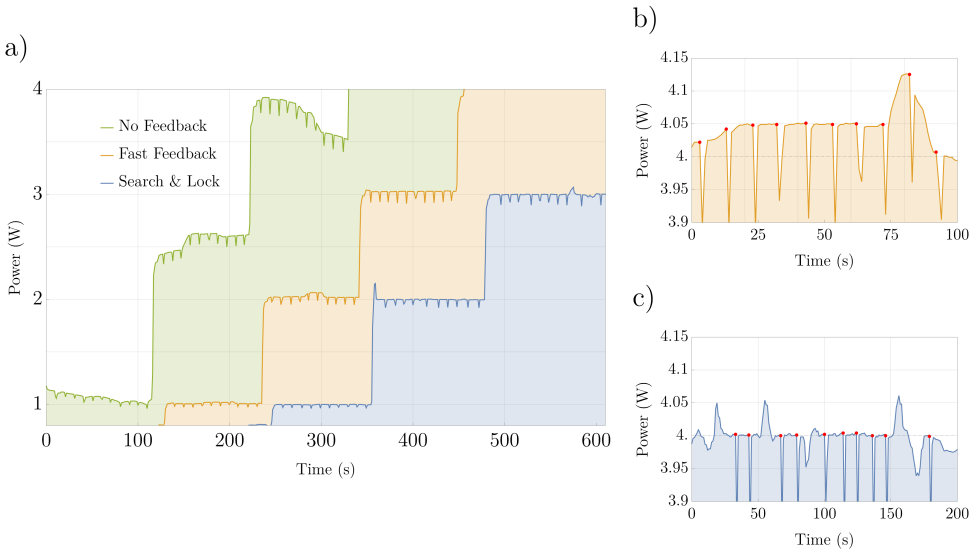

Figure 5.3: **Operation of the Search and Lock Algorithm. a)** A sample ablation sequence using improved methods of laser power stabilisation. There are three series of 10 shots at 1W, 2W & 3W, seen as dips in the measured power. The green trace shows the natural stability of the laser, which is greatly improved by application of fast feedback (yellow). The blue trace shows the dual-feedback method, showing both the short term stability of fast feedback and being closer to the intended power. **b)** 10 4W shots are performed with only fast-feedback, with the 4.05W mode dominating the sequence. Shot timings are highly periodic, since no laser re-locking was applied. **c)** The search and lock algorithm running for ten ablation pulses at 4W. Modal competition causes laser power to jump to 4.05W, but is successfully re-locked before each shot, shown in red.

The ablation beam is directed into the thermal power meter by an electronically controlled flip mirror.[8] Since power cannot be monitored during ablation, the final correction is made 0.8s before the laser is pulsed. This delay corresponds to the mechanical switching time between ablation and power-monitoring beam lines and sets the minimum period of stability required. Following a pulse, power monitoring is re-established and there is a waiting period to allow the substrate and meter to re-thermalise. For most ablation shots, the power variance over this period was low and re-locking of the laser was not required. However, if the deviation exceeds a defined tolerance, a search

[8]Thorlabs MFF101/M

and lock algorithm repeats until the desired error threshold is achieved. Accordingly, laser instability manifested itself as a longer delay between pulses, instead of a variation of their power.

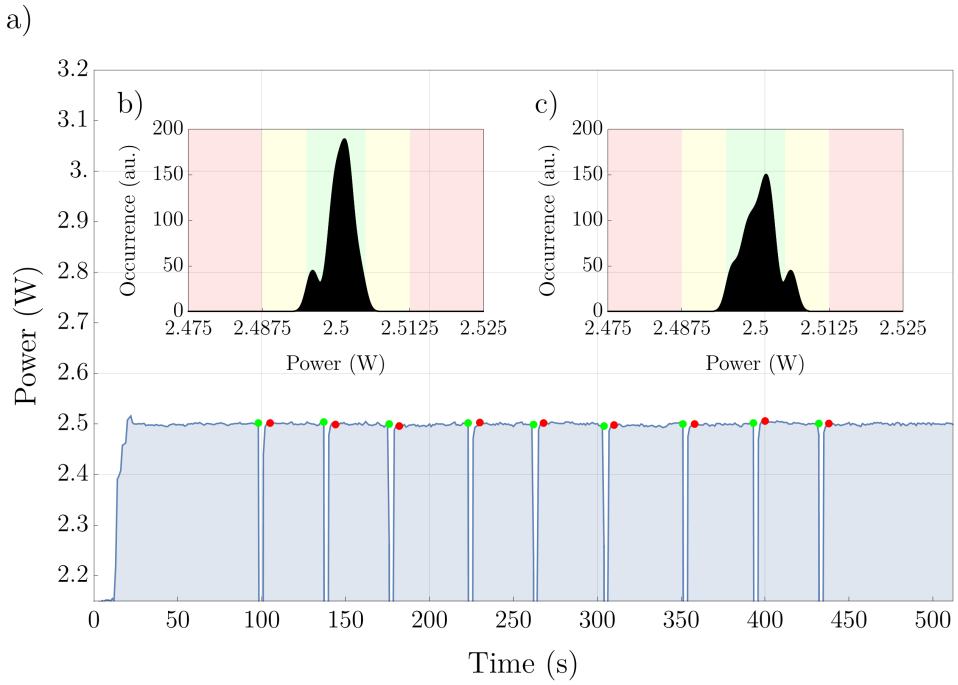

Figure 5.4: **Final Power Stability of the Laser. a)** A power trace for a standard ablation sequence is shown. Due to the electronic switching delay and power meter thermalisation time, there is an elongated delay between measurement of the power before each shot (green) and after (red). Each power dip corresponds to an ablation pulse being applied, where the duration is considerably shorter than the power meter integration time. **b)** The pre-shot power distribution is inset as a histogram, being compared to regions of $< 0.2\%$ power variation (green) and $< 0.5\%$ (yellow). **c)** The post-shot power distribution, also showing regions of $< 0.2\%$ and $< 0.5\%$ power variation.

The action of the re-locking algorithm is outlined fully in Figure 5.3. The intended sequence comprises three consecutive sets of ten identical shots, at powers of 1W, 2W and 3W respectively, representing the standard range used. Without feedback, the system had very poor performance, both in the average power and its standard deviation. Fast feedback improved consistency, but some mode hops persist over several shots and absolute power calibration slowly varies. Dual feedback rectifies calibration

errors, while introducing occasional delays between shots. This on average doubles the array fabrication time, for a defined tolerance of ±0.05W. The shot-to-shot noise of the dual feedback system is measured to have similar variance to only fast feedback, with the benefit of a consistent mean power. An example of the power spectrum of a representative ablation sequence is shown in Figure 5.4. It can be seen that power fluctuations are $< 0.5\%$, including the mechanical switching delay.

5.2.3 Pulse Shaping

The standard pulse length for ablation varies between 5 - 150 milliseconds, for which an accuracy in the microsecond regime is important. A diaphragm shutter[9] was initially tested as a method of beam shaping. The rise and fall time of the shutter was measured using a light gate arrangement, comprising the shutter, a laser diode and fast photodiode. The standard exposure accuracy of the shutter was ±1.8ms, which was improved to ±0.7ms by locating the shutter at the beam focus. This is slightly above desirable, but presented a secondary issue in that short-duration exposure patterns were dominated by the shutter transition profile.

Alternatively, the laser power can be quickly modulated via the RF controller, with a quoted accuracy of $< 150\mu$s. This performance could not be characterised directly, due to the lack of availability of MHz bandwidth photodiodes responsive at FIR wavelengths. While Mercury-Cadmium-Telluride (MCT) sensors are generally used at this wavelength, they are an order of magnitude more expensive than the industry standard silicon doped detectors used for visible wavelengths. As an alternative, the black-body radiation produced during the ablation of glass was used as a signature of laser irradiation. At very high powers a white glow is clearly visible from the heated surface, easily measurable by a Si detector. The results of this characterisation are shown in Figure 5.5. Due to a gradual heating of the surface, the FIR laser is not immediately detectable. However at the end of the pulse there is a sharp drop in irradiance, delayed

[9] Thorlabs SHB025

by only 200μs with respect to the control signal. Therefore, it was further assumed that laser gated pulses had a greater accuracy than those controlled by the shutter.

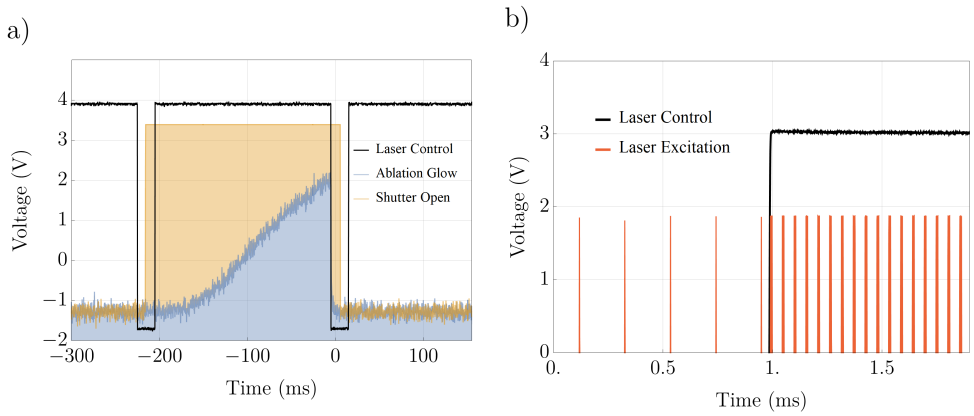

Figure 5.5: **Laser Gating Method for Pulse Shaping. a)** The control signal for the laser (black) is compared to an ablation glow measured on a photodiode (blue). Due to slow heating, the switch-on response of the laser cannot be determined, but the switch-off is clearly visible and as intended. This is compared to the period the shutter is open (yellow), timed to modulate when the laser is powered down. **b)** The first leading edge of the laser control signal is shown in black. As the laser output is activated the excitation scheme (red) changes from 5kHz tickle pulses to a 20kHz driving signal, pushing the laser over threshold.

To form each ablation pulse, the laser powers down, the beam is directed onto the substrate and a top-hat illumination profile is applied. Afterwards, the beam is directed away from the substrate and continuous lasing is re-established. In order to maximise the consistency of each exposure, the switch-off time of the laser should be kept minimal. Otherwise, the temperature of the gain medium and resonator length is likely to drift substantially, causing a modal shift once excitation is resumed. The transition speed of the flip mirror proved too slow to prevent this, so the shutter was installed as a secondary 'fast' modulated mirror. This directed the laser power into a graphite beam trap before each ablation pulse, requiring the laser to be powered down for only 20ms.

To ensure that an ablation pulse had the same peak power as during continuous op-

eration, tickle pulses are applied to the gain medium during the 'off' period. These excitation pulses had a duration of 1μs, supplied at a rate of 5kHz to ensure that the gas remains ionised, but slightly below lasing threshold. Ideally, this ensures the laser is ready for the rapid recommencement of output, as shown in Figure 5.5b). However, since there is no synchronisation between the tickle pulses and the laser pulse shaping, there is a random delay before each ablation shot. Empirically, this method of pulse generation was observed to greatly increase the prevalence of mode-hops, when compared to continuous laser operation.

5.2.4 Profile Smoothing

Despite significant efforts in improving the operational stability of our laser system and the accuracy of exposure duration, fluctuations remained visible in the features fabricated. This is commonplace in other ablation systems, where a large proportion of ablated fibre facets is simply discarded. However, in fabricating large arrays of features on the same substrate, this becomes impractical given the increased time and cost of replacement. The magnitude of these fluctuations was greater than one would anticipate, given our characterisation of pulse parameter accuracy and material sensitivity, suggesting a more obscure mechanism at play. It is likely that this resulted from our pulse shaping method, where the switch-on behaviour of the laser is naively assumed to be identical to that of continuous operation. Since we lacked measurement apparatus with sufficient temporal resolution to characterise this behaviour properly, it remained undiagnosed.

However, large arrays of excellent consistency were produced by the application of several pulses to each array site, thereby averaging out fluctuations. Since this was further improved by randomising the array coordinates between each pulse, it suggested that some slowly varying process was contributing to the error. A key difficulty was that the progressively changing profile of the glass surface reacted differently to consecutively applied pulses. As an example, if the first shot was especially powerful, creating a very

deep feature, it was almost impossible to recover a completely symmetric array. As such, weaker pulses were preferred, as a spurious mode would generally be recoverable. However, weaker pulses would generally require a larger number of repetitions to reach the desired geometry, requiring a longer fabrication period. When combined with the requirement to re-lock the laser power regularly, this period became limited by thermal drifts in substrate position.

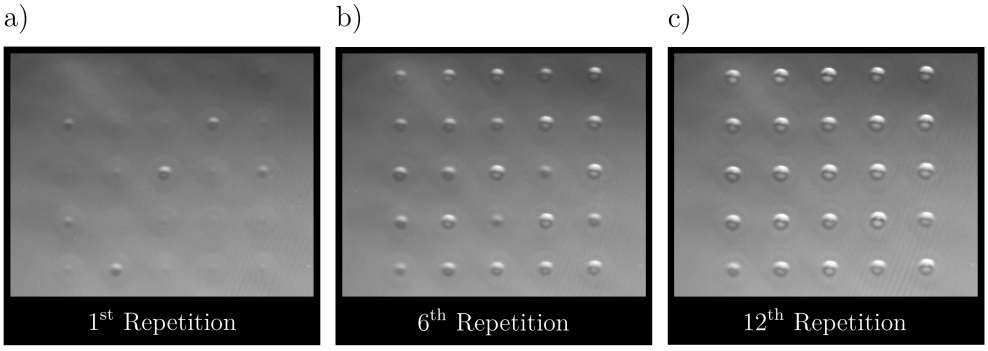

Figure 5.6: **Improving Regularity by Array Repetition.** The improvement of feature regularity can be seen on a 5×5 grid of features, each separated by $200\mu m$. The laser was re-locked after each shot, during which time the substrate was translated to the next ablation site. After each site was shot once, an image was taken and the pattern was repeated. **a)** Shows the initial array, **b)** shows the 6^{th} repetition and **c)** the 12^{th}.

A compromise came in applying staggered bursts of exposures to each site in the array. Laser power stabilisation was performed before each burst, normally comprising of approximately 5-10 exposures. After each burst, while the substrate was re-thermalsing, it would be translated to a randomly selected neighbouring site. Once every site had been addressed, the whole process was repeated, usually up to 15 times. An example of array profile smoothing is shown in Figure 5.6, which reliably produced symmetric patterns. However, this comes at detriment to the global applicability of our fabrication parameters, since they are now inherently tied to our systematic fluctuations.

5.2.5 Improved Design Schematic

Our analysis thus far has highlighted a number of technical improvements that could be made to the experimental apparatus. These are anticipated to greatly improve the reliability and reproducibility of ablation, at the cost of notable financial investment. In essence, these recommendations are generally the reapplication of laser control methods routinely applied at visible wavelengths, where active and passive optical devices are more readily available.

Our first requirement is to change the method of pulse shaping, such that it does not require modulation of the laser drive signal. An ideal method is an acoustic-optic modulator (AOM) based optical switch, which uses an AOM to modulate laser throughput. Therefore, by simply sweeping the drive voltage supplied to the AOM, we can create exposures with a temporal resolution much greater than currently possible. A more detailed description of this switching method is given in Chapter 6, where it is used to construct a cavity ringdown spectrometer.

Since the laser may now operate continuously, the prevalence of mode-hops will be greatly reduced. However, some power fluctuations will certainly persist. These could be reduced by monitoring a portion of the beam on an MCT photodiode and applying a correction signal to the AOM. Assuming there is no conflict in the required modulation frequencies, power correction and beam-shaping can be performed simultaneously, by simply mixing their AOM input signals.

To fully mitigate the presence of mode-hops in the laser output, a monochromator may be added to the beamline. If the laser then jumps between 10.53μm and 10.67μm lines, it is easily detectable and the target substrate is inherently protected. However, further studies have indicated that switching to a 9.3μm CO_2 laser is actually preferable for laser ablation, as it is absorbed more strongly by the glass [39]. This leads to a

reduction in the molten layer thickness, preventing the formation of convex features at the centre of the ablation site.

5.3 Substrate Positioning

So far, we have shaped the geometry of ablated features by careful specification of the laser exposure parameters. However, there is a limited range to these parameters, given the requirements of optical cavity mirrors. As a simple example, a short and intense laser pulse creates a very thin molten layer, generally proving insufficient to smooth the glass to optical quality. In contrast, a weak and sustained exposure will establish a predominantly molten layer, which under surface tension will form a convex feature. In order to achieve the enhanced diameters and controlled ellipticities recommended by our cavity design optimisation, these limitations must be overcome. This can be achieved by incorporating substrate translation into our ablation protocols. We will consider two established methods, substrate rotation and micro-patterned ablation. In this section we shall discuss the apparatus used to trial these simultaneously.

5.3.1 Mounting & Alignment

The majority of the degrees of freedom in substrate positioning are achieved using a 6-Axis *Nanomax*[10] stepper motor and closed loop piezoelectric stage. The Cartesian axes have a travel range of 4mm, limiting the maximum dimensions of a mirror array. The piezoelectric transducers have travel of 30μm, with a standard accuracy of $\pm$2nm. The three curvilinear axes are controlled by high-resolution manual micrometers[11] as they are only required during the initial alignment stage. A 360° bi-directional rotation stepper motor stage[12] is mounted onto the *Nanomax*, into which the target substrate is installed.

[10]Thorlabs MAX609L/M with DRV001 Stepper Motors.
[11]Thorlabs DRV3
[12]Thorlabs K10CR1/M

To allow for prototyping, a generalised substrate holder was attached to the rotational stage. This accepted smooth cylindrical elements of 25.4mm diameter, secured by Teflon setscrews. Adaptors were developed for mounting substrates smaller than this, down to 125μm diameter optical fibres. Light could pass through the rotational mount, allowing for simple optical alignment when a pair of irises were attached.

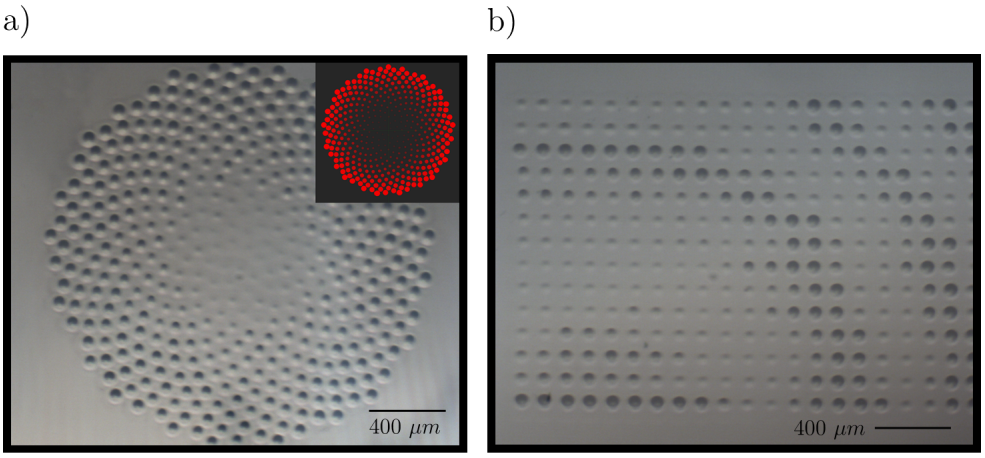

Figure 5.7: **Substrate Positioning System. a)** An optical microscope image demonstrates that precise mirror patterns can be constructed by autonomous control of the stepper motors, here showing $\sim$ 400 features in a phyllotactic arrangement. Inset is the ablation fabrication coordinates uploaded for manufacture. **b)** A regular array of single shot features with equal power, duration and spacing was fabricated. The combination of arcuate motion and cross-talk due to the multi-axis flexure state is shown, demonstrating oscillation in substrate depth when translated perpendicularly to the beam.

To ensure alignment between the ablation beam and target substrate, a visible 640nm laser diode beam was overlapped the CO_2 light. This wavelength was selected as it is partially reflected by the common optical material zinc selenide, whilst mid-infrared wavelengths are transmitted. As such, a simple zinc selenide laser window acted as a longpass dichroic filter.

The laser diode beam profile exhibited significant ellipticity, which made it difficult to check lens alignment for astigmatism. This could have been rectified by replacing it

with a Helium-Neon laser, which operates at a similar wavelength with considerably improved modal quality. However, the greatest signature of system alignment came in the fabrication of feature arrays. If the ablation beam was not addressing the target substrate perpendicularly, then feature depth would gradually shift across the array. By estimating the distance from focus of each feature, one could adjust alignment proportionally. A complex array pattern is shown in Figure 5.7a), demonstrating good system alignment.

5.3.2 Translational Cross-Talk

One can note that small deviations are visible in this array pattern, which have largely been attributed to the laser instability issues discussed previously. However, it was later detected that there was a slight correlation between error size and location on the substrate. To examine this further, arrays of very close feature spacing were fabricated, with a site separation of only 100μm. This was previously avoided for standard array fabrication, as notable overlap of each feature occurs. However, this highlighted an oscillating pattern, where depth varied in waves emanating from a fixed point on the substrate. An example of this pattern is shown in Figure 5.7b).

The flatness of the substrates was measured by coherence-scanning interferometry, but showed deviations significantly smaller than could cause the observed pattern. Instead, this behaviour appeared to be caused by cross-talk in the *Nanomax* translation stage: as the substrate was moved to the position of the next site, some oscillation in the distance of the substrate to the focusing lens would occur.[13] This derived from the stage operating using a flexure mechanism, rather than having independent movement axes. This platform was selected as it allowed for 12 movement actuators in a very compact profile, but this clearly came at expense to their translational accuracy.

These flexure patterns seemed to have history dependence, so were greatly reduced in

[13]The manual of the *NanoMax* stage referred to this as arcuate motion.

magnitude by randomising the order of feature fabrication. This was integrated into the profile smoothing methods of Section 5.2.4. Ideally, there would be an independent system developed to monitor the location of the translation stage, allowing suitable corrections to be made.

5.3.3 Rotational Smoothing

In the array of ablation features presented previously, one could note identical ellipticities, elongated along the horizontal axis. As the expected cause, the profile of the laser beam was characterised using a knife-edge method and indeed presented a near identical profile. In the first instance, we require perfectly spherical cavity mirrors to prevent birefringence. The simplest way to rectify this issue was the rotational smoothing technique trailed in our research visit to the Keller group [40] and recreated in our apparatus. This applies a large number of ablation pulses as the substrate is rotated, averaging out any inherent beam ellipticity. Without the use of any beam shaping optics, features with eccentricity $< 2\%$ could be readily fabricated.

A key difficulty in this method is that it is exceptionally challenging to ensure that the axis of substrate rotation is perfectly aligned with that of the ablation beam propagation. Therefore after each rotation step, Cartesian translations must be performed to bring the feature site back into alignment. Given that up to 50 shots would generally be applied to each feature site, this was unfeasible to perform manually. Instead, an algorithm was developed which could recognise the position of the substrate by optical imaging and automatically apply the required corrections. This proved effective for fibre production, as the illuminated edge of the fiber presented an easily resolvable indicator of position. However, the increased size and irregularity of macroscopic substrates proved challenging for the development of a universal object recognition protocol.

5.3.4 Micro-Patterned Ablation

A further limitation of rotational smoothing is that ellipticity is simply removed, rather than being controlled. As discussed in Section 4.6.5, ellipticity actually presents an interesting degree of freedom in polarisation dependent photon production schemes. In order to tailor it into our features we performed micro-patterned ablation, or 'laser-dot milling' [47]. This fabricated a feature using many laser exposures, each separated by small piezoelectric motions.

Our standard pattern was to follow sets of concentric ellipses, gradually migrating towards the centre of the feature. The final laser exposure was at the feature centre, given the intended location of the cavity mode. In order to create spherical features from the elliptical beam, these rings were given an eccentricity equal to that of the laser profile, but with a perpendicularly aligned major axis. An example micro-pattern is given as Figure 5.8. In general, feature eccentricity could be enhanced or mitigated by stretching of the milling pattern.

A second advantage to this method is that the feature diameter may be enhanced. In general, this allowed us to expand the standard 40μm feature diameter of a single site up to 100μm for a patterned site. This was limited both by the maximum translation range of the piezoelectric stage and by the inherent ellipticity of the beam: when creating spherical features, some trade-off between ellipticity cancellation and diameter enhancement was made. Since this is a critical characterisation of our ablation apparatus, it is considered in further detail in Section 6.3.2, in the context of associated spectral analysis.

5.4 Experimental Control

Fabricating an array of individually micro-patterned features requires several hundred substrate translations, frequent laser re-stabilisation and surface inspections. A crucial

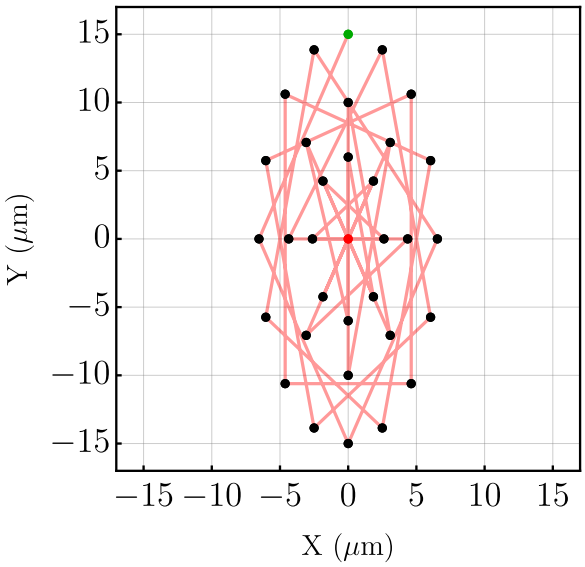

Figure 5.8: **Micro-Patterned Ablation.** An example micro-pattern is shown, where each point represents a laser exposure position on the substrate. This pattern consists of three ellipses, comprising 16 exposure positions in the outermost ellipse and 8 exposures in each of the two inner ellipses. The first exposure is at the top of the pattern, shown in green. The final exposure is at the feature centre, shown in red. The route between consecutive exposures is shown as a red line. Where possible, ablating two neighbouring sites in a row is avoided, to ensure creation of an even profile. The ellipticity of this pattern was sufficient to negate the ellipticity in the ablation laser, giving spherical mirror profiles.

element of the fabrication platform was therefore the centralised control of all devices electronically. Software was written in Mathematica, which allowed the user to specify a coordinate for each feature, including their location, pattern, laser power and pulse duration. A full set of coordinates was then compiled into an instruction set file, which could be uploaded to a graphical user interface built in MATLAB. From this interface, every element of the apparatus could be controlled manually, or when an instruction set was loaded, an array could be fabricated and evaluated autonomously. This facility was naturally reliant upon a large set of interconnected electronics, much of which was constructed specifically for this experimentation. We shall now outline the key elements.

Laser power is governed by the duty cycle of a 20kHz signal, driving the excitation of

the gain medium. This signal was generated by a Synrad UC2000 controller, which linearly maps an input 0-10V DC control signal to 0-100% duty. This DC control signal was generated using a low noise digital to analogue conversion (DAC) chip,[14] mounted on an Arduino Due. The Arduino was used to support USB connection to the PC, passing instructions to the DAC chip by a 12 bit parallel connection. Since the Arduino and PC have asynchronous clocks, a dual stage latch was used to time data transfer.

For beam shaping, laser output was rapidly modulated, controlled by a transistor logic (TTL) input signal to the UC2000 laser controller. Defined as the 'gating' input, this was independent to the duty cycle specification. This signal was synthesised using an Arduino Uno to support the USB interface and an Arduino Due to manage TTL signal output. Sequence timing was passed between the Arduino boards using an 8-bit parallel connection, allowing pulses to have a 0-255ms duration, specified at increments of 1ms. The accuracy of pulse duration was limited by the Arduino Due clock frequency of 84MHz. For stability, the laser was kept running between ablation sequences, during which all Arduino outputs were suppressed to allow for reprogramming. Therefore, a XOR gate was placed on the gating input, so that logic high switched the laser off.

Since shutter operation must be synchronised to the laser gating signal, they were supplied by the same digital sequencer. The shutter had an integrated control circuit, opening on a logic high signal output directly from the Arduino Due. The delay between the laser and shutter operation was accounted for, calibrated using a standard light gate measurement.

Operation of the *NanoMax* stepper motors was performed using a Thorlabs BSC203 controller unit. The drive and feedback of the piezoelectric transducers were operated equivalently, using a Thorlabs MT3101 controller. The rotational stage had an inte-

[14] Analog Devices AD669

grated controller, connected directly to the PC. However, the wire to the rotational stage was seen to transmit vibrations into the system, so was only connected during rotational smoothing. Otherwise, its manual actuation was used to align the axis of feature ellipticity to that of the rectangular pyramidal mirror facets discussed later. Each of the three flip mirrors had integrated controllers, connected directly to the PC. All devices utilised a USB 2.0/3.0 connection and were programmed using ActiveX protocols and Thorlabs 32-bit APT drivers.

In the following section, we shall consider in detail surface reconstruction methods. This utilised a tunable Michelson interferometer, whose piezo[15] was driven by a high voltage amplifier[16] and controlled by a second Arduino-DAC unit. Inspection was performed using a CMOS sensor,[17] connected to the PC by a USB 3.0 connection and programmed by the MATLAB Image Acquisition Toolbox.

5.5 Surface Reconstruction

The final element of an ablation platform is an integrated facility for the inspection of substrates. System alignment, calibration and elementary feature characterisation can all be performed by simple optical imaging. However, a more informative method of surface reconstruction must be sought in order to overcome the unpredictable nature of thermal processes: since the pulse sequence which corresponds to a desired mirror geometry cannot be predicted analytically, it must be searched for empirically. To create bespoke mirrors in a time-efficient manner, we therefore require a rapid, non-destructive and in situ method of topographical reconstruction.

There are several properties of interest in potential cavity mirrors. At the forefront is the geometrical parameters of the depression, defined by its curvature, diameter and el-

[15]Thorlabs PK4GYP2
[16]Thorlabs MDT694B
[17]Imaging Source DFK 23UP1300

lipticity. In order to anticipate the optical properties of the mirror post-coating, we also need to evaluate surface roughness and detect impurities. No single technique is able to perform a comprehensive characterisation, so several have been employed. Geometrical characterisation is performed by phase and coherence scanning interferometery [135], integrated into the standard ablation process. External to this, sub-micron topography is characterised by atomic force microscopy [136] and substrate contamination is monitored by dark-field microscopy. In this section we shall discuss the design and implementation of these systems.

5.5.1 Phase Scanning Interferometry

Phase scanning interferometry (PSI) is a method of surface profilometry, generally performed using broadband illumination [135]. It uses an optical arrangement akin to a Michelson interferometer, with a reference and target substrate being located on neighbouring arms of a beam splitter. The light reflected from their surfaces is mixed on a third arm, creating an interference pattern. Inhomogeneity in this interference pattern can be interpreted as a path length difference between the target and reference arms, essentially denoting variance in the substrates. For a perfectly aligned and flat reference substrate, interference lines can be directly interpreted as topographical contours of the ablated feature, allowing for immediate profile characterisation.

Technical Outline

The optics used to construct this interferometer are naturally incompatible with the ablation beam, being entirely opaque towards mid-infrared wavelengths. Therefore, two electronic flip mirrors were used to introduce the beam-splitter and illumination mirror to the system when required. This mirror was illuminated by a halogen bulb,[18] which was coupled into a fibre bundle to isolate the generated heat and fan vibrations. The reference substrate[19] was refractive index matched to the ablation glass, with $\frac{\lambda}{20}$

[18] Thorlabs OSL2
[19] Thorlabs SSM0510B

flatness and $\sigma_{rms} = 1$Å surface roughness.

The most challenging aspect of PSI was to ensure the alignment of the reference substrate. Any surface gradient would cause the fringe spacing to rapidly compress below the resolution limit of the imaging arm. Then, given the short coherence length of the illumination source, the distance of the substrate to the beam splitter had to match that of the target substrate to micron accuracy. To achieve this, the reference substrate was attached to a piezo stack, in turn attached to a kinematic tip/tilt mount. To give a greater range of potential path lengths, this was mounted onto a single axis manual translation stage[20] with 6mm travel.

The standard distance of each substrate from the beam splitter was determined by the working distance of the microscope. A zoom telescope[21] was used with a relatively long working distance of 34mm, to allow sufficient space for retracting the beamsplitter during ablation. The maximum magnification of this telescope was 20x, which when combined with the 2.5" colour CMOS sensor, gave a field of view of 330μm $\times$ 220μm. At the minimum magnification the field of view expanded to 1.1mm $\times$ 0.8mm, allowing for the standard optical inspection of large substrate areas.

The use of a colour sensor allowed one to compare different spectral bands of the interference pattern. A shift in the relative location of each band maxima would denote a change in the gradient direction of surface topography. Since convex features are known to form if the ablation exposure creates an overly thick molten layer, this proved a useful diagnostic tool. However, for standard reconstruction only the red channel was interpreted, as it more closely matched the illumination bulb temperature and gave the strongest signal.

[20] Thorlabs MS1S/M
[21] Navitar 1-50487 with attachments.

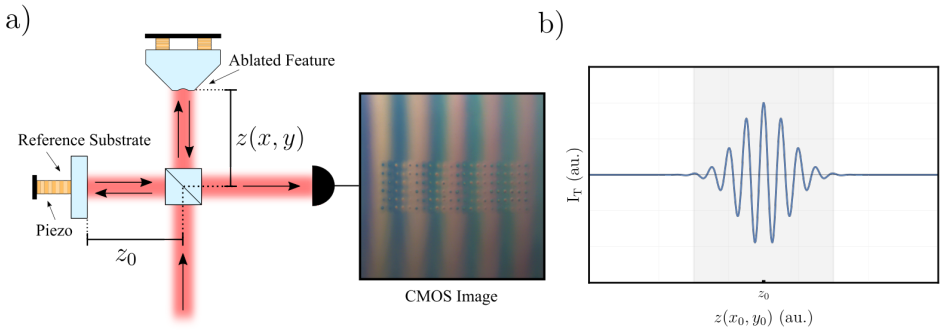

Figure 5.9: **Interferometric Reconstruction.** **a)** A schematic of the Michelson interferometer used in surface reconstruction is shown. An incoherent input beam is used to illuminate a flat reference mirror and ablated feature. Their reflected light is mixed on a beam-splitter and resulting interference pattern recorded on a CMOS sensor. A sample static inteferogram is given, showing an array of features on a flat, horizontally tilted substrate. When the distance from the beam-splitter to the ablation feature, $z(x,y)$, matches the distance to the reference substrate, z_0, an interference maxima is observed in all colour channels. As one moves further right in the image, $z(x,y) > z_0$ and the spectral bands separate. The spacing between each colour maxima is $\lambda_0/2$, where λ_0 is a combination of the central wavelength of illumination source and CMOS colour channel filter. **b)** In coherence scanning interferometry we scan z using a piezo, resolving each pixel of the interferogram, (x_i, y_i), individually. An example intensity profile is shown, denoting interference only within a coherence envelope. The maxima of this envelope corresponds to $z(x_i, y_i) = z_0$, giving a measure of surface topography when all pixels are fitted. The shaded region corresponds to $z = z_0 \pm L_c$, where L_c is the coherence length of the selected colour channel.

System Calibration

The relationship between the observed interference pattern, $I(x,y)$, and the height map of the target substrate, $z(x,y)$, is given by:

$$I(x,y) = I_a + I_b\, e^{-4\left(\frac{z(x,y)-z_0}{L_c}\right)^2} \cos\left(\frac{4\pi}{\lambda_0}(z(x,y)-z_0) - \phi_A\right). \tag{5.1}$$

Where $I_{a,b}$ are fixed constants and z_0 is the reference substrate position with respect to the beam-splitter. For simplicity we have neglected to consider the geometry of illumination in the interferometer, characterising it only by an effective wavelength λ_0.

All remaining phase shifting effects are combined in ϕ_A. L_c encodes the coherence length of the illumination source, given by $L_c = \frac{c}{\pi\Delta\nu}$, where $\Delta\nu$ is the approximate 1/e spectral bandwidth of the illumination source. A schematic of the interferometer is given as Figure 5.9, further describing each parameter.

In phase scanning interferometry we consider only a static interferogram, corresponding to a fixed value of z_0. To convert the observed fringe pattern into a calibrated surface map, the value of λ_0 must be accurately determined. Since this is a complex product of the halogen bulb output, coloured filter and dispersion by optical elements, λ_0 was determined in situ, by a modified method of interferometric reconstruction. This made use of the coherence envelope, given by the exponential term in the above expression. The peak of this envelope is observed as the position of the brightest fringe for each colour channel, easily recognisable as a white band. Since this corresponds to $z(x, y) = z_0$, if z_0 is precisely scanned, then substrate geometry can be derived by the surface migration of the coherence maxima. This forms the fundamental principle of coherence scanning interferometry, which shall be considered in more detail in the following section.

To determine λ_0, we performed coherence scanning interferometry on a flat aluminium mirror, tilted at a small angle, θ, in the horizontal axis.[22] The static phase pattern is vertical coloured bands, whose spacing was defined by θ and λ_0. Using the piezo stage the aluminium plate was retracted 10μm, with interferograms recorded at 0.1μm intervals. By recording the position of the white band at each frame, θ could be determined. Then, using the static phase image, the effective wavelength of each colour channel could be derived. For the red channel this was $\lambda = 638$nm, meaning that each interference fringe extrema represented a 152nm change in feature depth.

[22] With respect to the reference substrate, which was assumed to be perfectly perpendicular to the illuminating beam.

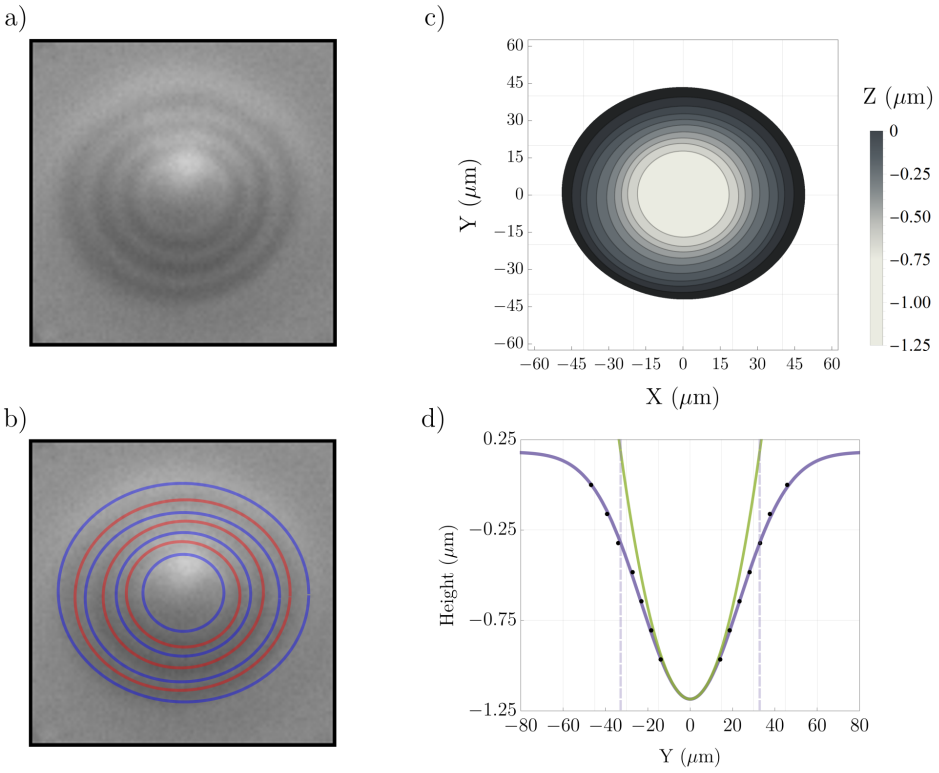

Figure 5.10: **Phase Reconstruction of Features. a)** A red-channel image of a standard single-shot ablation feature. **b)** Output of the ellipse fitting algorithm, fitting contours to maxima (blue) and minima (red). **c)** These contours can be calibrated to depths using the effective wavelength of the light source. **d)** A fitted Gaussian (purple) is shown for a central cross section in the Y axis of this feature. The boundaries of its 1/e waist is shown as vertical dashed lines. A corresponding spherical profile is shown in green, fitted to the central region of the Gaussian profile.

A calibrated interferogram of a single-shot ablated feature is shown in Figure 5.10. It can be observed that the depression has a highly Gaussian profile, closely resembling the intensity distribution of the ablation laser. Only the central portion of this feature can be considered as near-spherical, from which we define the geometrical parameters of the mirror. However, this approximation is subjective, with the boundary of the 'spherical' section having no formal definition. To ensure that our fitting approach is representative of the optical properties of each feature, they are compared to geometric measurements made using cavity-based spectroscopy. This is performed in Section 6.3.2, where we make a detailed evaluation of the accuracy in our phase scanning

reconstruction method.

Evaluation

A key difficulty with phase scanning interferometry came from the signal to noise ratio. The substrates used to support feature arrays consisted of two large polished facets, where reflections from the back surface would mix with that of the ablated side. This reduced fringe contrast significantly, making it difficult to profile deep ablation features. The illumination beam was apertured to reduce stray illumination, but this was limited due to the compact footprint of the interferometer. Alternatively, the spectral density of a chosen colour channel could be increased by using a high power LED. While the gradient direction of a feature would no longer be represented, a greater illumination with finite coherence length would be achieved. Fringe contrast did not prove an issue for producing tapered mirrors, owing to less stray reflection from their rear surfaces.

5.5.2 Coherence Scanning Interferometery

In order to overcome the resolution limitation of phase scanning interferometry, the finite coherence length of the illumination source can be utilised. In contrast to a conventional Michelson interferometer using monochromatic light, white light will only generate an interference pattern over path length differences of up to $\approx 10\mu$m as a result of its lower spectral coherence. Over this path length, a Gaussian envelope is superimposed onto the interference pattern, the centre of which denotes no difference in path length. During this scan, each pixel of the CMOS array essentially becomes an independent interferometer, allowing a far more detailed reconstruction of the surface topography. As such, this method is able to determine the topography of highly scattering surfaces, with features un-resolvable by a static interfeometric image. A representation of the interference pattern measured for each pixel is given as Figure 5.9b).

However, in requiring multiple images for surface reconstruction, the raw data set becomes much larger. For a full reconstruction of the surface, generally each of the

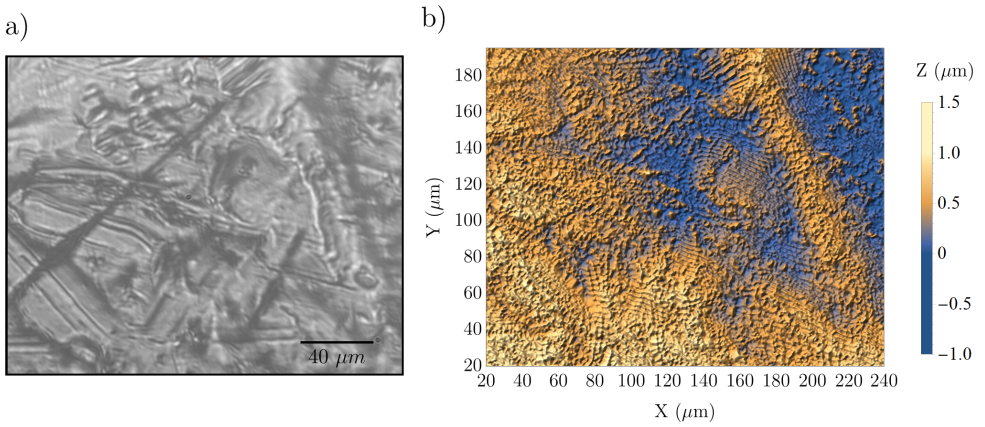

Figure 5.11: **Noise Floor of Coherence Scanning Interferometry. a)** An optical microscope image of a tarnished coin, where surface topography is difficult to determine. **b)** A coherence scanning reconstruction of the coin surface, showing sub-micron topography. Systematic noise can be seen as highly periodic oscillations across the surface of the coin. Even with the coin uncleaned, 85% of the pixels yielded high-quality fits. Missing data points are accounted for by a 5 pixel Gaussian filter.

$\approx$ 1,000,000 pixels must be independently fitted. In practice, downscaling of the image is performed to average out noise and reduce processing time. Numerical methods can also be employed, such as Hilbert transformations, to further improve the efficiency of fitting methods [135]. In this study, the surface reconstruction is performed using numerical fitting procedures native to *Mathematica*.

The resolution of this method is defined by the fitting accuracy of the coherence envelope. In practice, this is limited by the vibrational noise of the system. Since the piezoelectric stage operates with closed-loop feedback for accurate translation, it moves comparatively slowly between imaging points. Therefore, the interferometer was susceptible to environmental fluctuation and displayed a standard noise floor of 0.4μm. This was demonstrated in the interferometric reconstruction of a tarnished coin, given as Figure 5.11. This technique could certainly have been improved by enhancing the stability of optical mounting, albeit at detriment to its integration with the ablation beamline.

5.5.3 Atomic Force Microscopy

Surface imperfections that approach an atomic scale in dimension cannot be resolved by white light interferometry. However, surface roughness on the scale of $\sigma_{rms} < 1$nm is the key cause of scattering losses [44]. In order to determine the surface roughness of ablated features, an atomic force microscope is used [136].

An atomic force microscope is historically derived from a stylus profiler. A stylus profiler would drag a lever across a sample surface, using the deflection force experienced to determine changes in sample height. However, this would prove highly destructive to the surface in question, as well as having a poor resolution and noise floor. In contrast, an atomic force microscope oscillates a cantilever above the surface, whose frequency will change with proximity to the sample, as a result of their mutual electrostatic interaction. A feedback lock is applied to this oscillation frequency, whose error signal is interpreted as a signature of sample topography. There are two general modes of operation for an AFM, intermittent contact (tapping) mode and non-contact modes. Non-contact mode utilises the attractive electrostatic interaction found very close to the target surface, but therefore requires a very stable environment to maintain this separation. Instead, an intermittent contact mode was used for this work, which operates with a mainly repulsive interaction. However, the cantilever will briefly touch each measurement site, making the process more invasive.

The device used is a Nanosurf Nanite, operating with a C3000 controller and software suite. Tap190Al-G probes were installed, with a quoted noise floor of 10pm. In order to reduce environmental noise, the system was located in an acoustically isolated chamber, floated on an optical table. The maximum scan area is 75μm $\times$ 75μm. For the resolution of surface roughness, areas of 2.5μm$\times$2.5μm were more commonly considered.

Since intermittent contact mode is an invasive measurement, it was not integrated

into the standard mirror fabrication sequence. Measurements have the potential to introduce contamination to the mirrors, either due to direct contact with the cantilever or through the additional time they are exposed to unfiltered air. Instead, example substrates were scanned from each fabrication regime, to ensure that the surface tension of the molten layer was correctly smoothing the mirror surfaces towards atomic quality. A demonstration of this is given in the following section.

5.6 Pyramidal Mirror Manufacture

5.6.1 Review of Developed Platform

We are now in a position to review our developed platform. Our favoured approach to tailor feature geometry was to use micro-patterned ablation. This allowed the expansion of feature diameter from 40μm to 100μm, reaching the state of the art reported in other studies [47]. This permits the fabrication of a millimetre long cavity without clipping loss, a notable advantage for the facilitation of an intra-cavity dipole or ion trap. By careful specification of pattern eccentricity, the inherent ellipticity of ablation features could be tailored or mostly negated. In reducing feature ellipticity, the tested method of substrate rotation may be favourable, but otherwise provides less flexibility.

Two methods of surface profilometry were developed, which allowed the empirical connection of pulse sequences to feature geometry. A large range of feature curvatures were created, from $R_{oc} = 200\mu$m to $R_{oc} = 1750$mm, commensurate with the standard range of ablation outlined in Section 4.3.4. Phase scanning interferometry was favoured in this manufacture, given its relative simplicity and speed. Surface roughness was then estimated by atomic force microscopy, providing an estimation of optical losses when reflective coatings are applied to the substrates.

Our system of electronic control allowed the ablation and reconstruction process to proceed with minimal user input. This allowed the creation of arrays with several hundred

features. Ablation pulses were synthesised with sufficient accuracy in duration and intensity to produce consistent feature geometries, when repetition was performed to smooth noise. This reliance on repetition could be reduced in a future implementation, via the introduction of an acousto-optic modulator for pulse shaping.

As a prototypical batch, we explore the creation of large diameter mirrors, mirrors of tailored ellipticity and small arrays of features. Our choice of substrate must allow open access to the cavity mode, but is otherwise highly flexible in design. In the following sections, we motivate the design of these mirrors and detail the process of their fabrication and coating.

5.6.2 An Alternative to Fibre Cavities

In our discussion of cavity design optimisation, we observed that the standard fibre-tip facet diameter of $\sigma_s = 125\mu$m would meet our requirement for optical access to the cavity mode. This gave us greater freedom to tune cavity length, based on the optimisation of photon output efficiency and cavity linewidth. However, there are a range of experimental implications in the use of fibre-tip mirrors, which we are less advantageous.

To preserve the spatial coherence of cavity generated photons, the transmissive mirror must be placed on a single mode fibre-tip. Out-coupling then relies upon the geometrical overlap of the cavity mode with that of the fibrecore. This dependence on cavity geometry can cause mode-matching to be a dominant aspect of cavity design optimisation, or a notable limitation to photon output efficiency. A developed solution is to splice gradient index fibres onto the single mode fibre, effectively introducing a mode-matching lens to the cavity [46]. However, both fibre-tips and spliced fibres are known to be exceptionally fragile, making their handling and manipulation into a resonating cavity more challenging.

Fibre-tip mirrors are also recognised as having an inherently lower vibrational stability

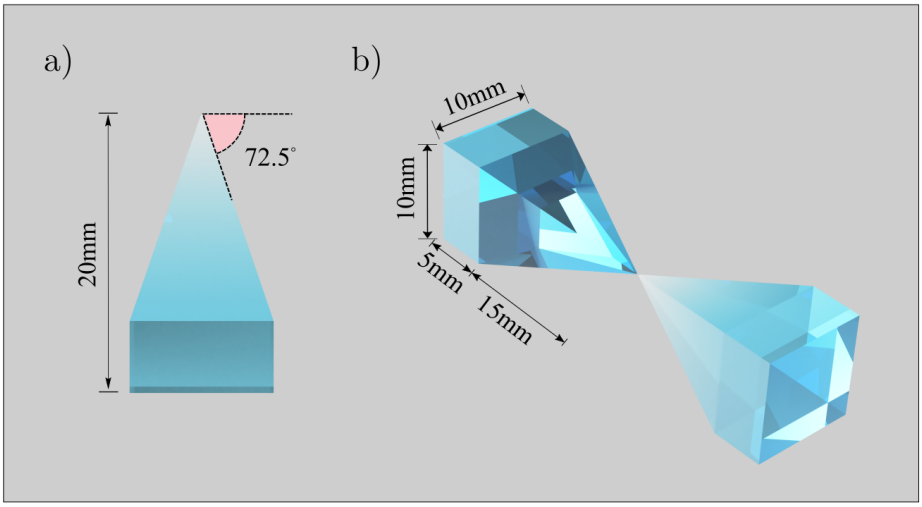

Figure 5.12: **Schematic of Pyramidal Mirrors.** **a)** A basic schematic of a pyramidal substrate is shown, highlighting its exterior dimensions and taper angle. **b)** Our expectation is to form a cavity using a pair of pyramidal substrates, as illustrated. This allows for a facet diameter similar to an optical fibre-tip, but a mirror that can be mounted as a macroscopic substrate.

than macroscopic cavities, due to their small mass and mechanical cantilever modes of oscillation [45]. Their mounting structures must therefore be carefully developed to suppress vibration and may restrict the bandwidth of cavity length modulation that can be applied [48, 49]. Also, the transport of heat within fibre-tips is restricted, which is absorbed from laser illumination of the cavity. The generation of an optical signal for cavity locking may therefore cause instability, by thermally expanding the length of the cavity as it approaches resonance [50].

To avoid the technical issues presented above, we proposed the development of a cavity mirror with the facet geometry of an optical fibre, but with the experimental practicality of macroscopic optics. Our chosen geometry was a square-based pyramid, with a flattened tip of area 140μm $\times$140μm and larger. It is now our intention to demonstrate the same optical quality and geometrical freedom as fibre-tip ablated mirrors, but with enhanced mode-matching, vibrational stability, thermal stability and ease of cavity construction. However, the use of pyramidal mirror substrates naturally intro-

duced a novel set of engineering challenges. We shall now outline our refined fabrication process, addressing each of these in detail.

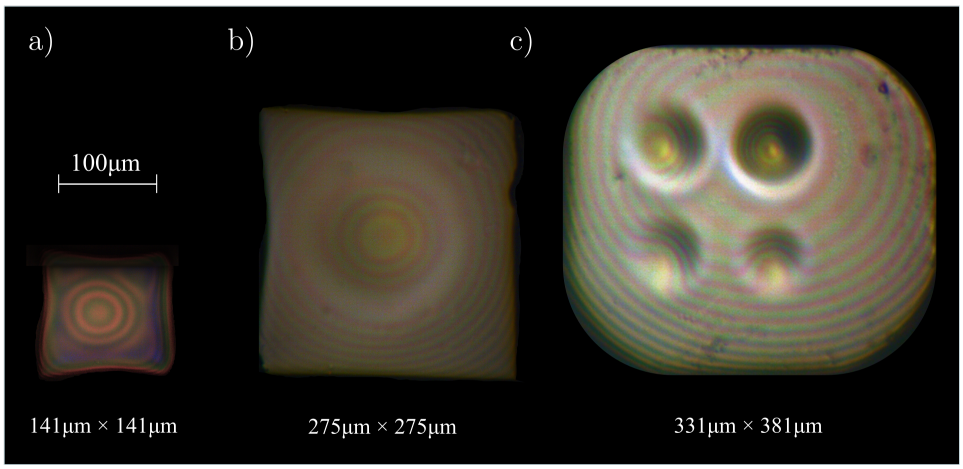

Figure 5.13: **Effect of Facet Size on Feature Geometry.** We show a phase-interferogram for three pyramidal mirrors, representing distinct facet sizes. All images are on the same scale, labelled by facet length × height in μm. **a)** The smallest facet produced was $141\mu\text{m}\times141\mu\text{m}$, intended to emulate a standard fibre-tip diameter. However, feature geometry was subject to notable warping from the edges of the substrate, reducing its effective diameter and increasing ellipticity. **b)** When the substrate was expanded to $275\mu\text{m}\times275\mu\text{m}$, larger mirror diameters were observed with more consistent ellipticity. **c)** An example array of features is shown, on a large facet diameter. At this size scale edge effects could be largely ignored, allowing low-ellipticity features to be fabricated using a suitable micro-pattern. However, the spacing between features was limited to $> 100\mu$m, to prevent warping via their overlap. This array uses the minimum feature spacing.

5.6.3 Glass Shaping

The most important physical dimensions of the pyramidal mirrors were their taper angle and facet size. The taper angle must allow internal expansion of the cavity mode, without it reaching the edge of the substrate. Since we intended to trial a range of mode geometries, a conservative angle of 72.5° was selected, as indicated in Figure 5.12. However, there was no single facet size that would facilitate all of our intended experimentation. For single features, the facet size would ideally be equal to the feature diameter, $\sigma_s = \sigma_m$, following our discussions in Chapter 4. However, this is impractical

when performing ablation, as the molten layer may extend over the substrate edges, warping feature geometry. For multi-feature resonators a considerably larger facet is required, given by the dimensions of the array.

UV-fused quartz was used as the base substrate, given its high optical quality and transmission at NIR wavelengths. Bulk substrate shaping was performed by a commercial optics manufacturer, with flattening of the pyramidal tip performed within our own laboratory. This was performed using a random orbital fibre-tip polisher[23] and custom mounting fixture. Several stages of polishing was applied to each tip, with abrasive grit size reduced down to 0.02μm. The polishing time of each stage was carefully controlled, to create the facet diameter of choice. However, owing to some error in the shaping of the base substrate by the commercial supplier, some facets proved slightly rectangular.

5.6.4 Ablation Process

The process of micro-patterned ablation on the pyramidal mirrors proceeded as for macroscopic substrates, using phase-scanning interferometry for surface reconstruction. However, the geometry of produced features was distinct, owing to the modified heat transport properties of the tapered substrates. The pyramidal tips were naturally more sensitive to applied laser exposures, creating larger depressions for a given pulse sequence. Therefore, the geometrically calibrated pulse sequences developed using macroscopic substrates could not be applied here. While there was clearly some proportionality between the pulse sequences applied to macroscopic and tapered substrates, it required several further iterations to produce a mirror of choice.[24]

A presented issue was that the rectangular facet did not match the spherical profile of the intended mirror. Therefore, the surface tension applied to the molten glass was

[23] KrellTech SpecPro

[24] It should be noted here that it is more trivial to cleave a new fibre facet than it is to shape a new pyramid. However, after a non-desirable feature geometry is applied to a pyramidal mirror, it may be polished away.

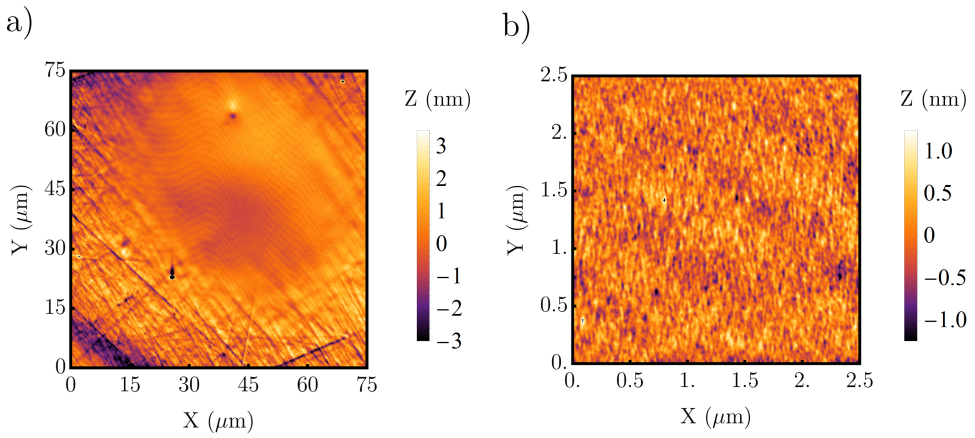

Figure 5.14: **Surface Healing by Laser Ablation. a)** The topography of an ablated pyramidal mirror is shown, measured using an AFM. Striations and surface damage from tip polishing can be seen in the periphery of the image. However, towards the site of ablation these surface deviations have been smoothed by thermal annealing. **b)** A $2.5\mu m^2$ region at the centre of the ablation feature is shown, from which surface roughness can be determined. This was measured to be 3.3Å, with a polynomial fit applied to remove feature curvature.

non-isotropic, stretching the ablated features. This effect can be seen in Figure 5.13a) and b), where two features of similar pulse pattern were applied to different facet geometries. On the 141μm facet, feature ellipticity has been enhanced in the direction where it approaches the facet edge. However, when expanded to $\approx 275\mu$m in b), significantly less feature warping has been induced.

To approach the diameter of fibre-tips, $\sigma_s = 125\mu$m, it may be advisable to use tapered substrates with cylindrical symmetry. However, edge effects are also present in fibre-tips, where the sharp facet edges are highly susceptible to melting [115]. The rectangular geometry proved ideal for the production of feature arrays, comprising a square lattice of spherical depressions. An example of such an array is given as Figure 5.13c), comprising 4 micro-patterned features. Since we are constrained to a large facet size by the spacing of the features, issues of heat transport were less prevalent here.

The final concern was the quality of the optical surface produced by standard polishing.

By using abrasive paper, the pre-ablation surface roughness of the pyramidal facets was several orders of magnitude greater than required for low-loss cavity mirrors. Therefore, they were fully reliant on the thermal annealing effect associated with ablation to reduce scattering losses. The effect of 'healing' surface imperfections in the molten glass layer was examined under the AFM. A key representation of this action is given as Figure 5.14, where the ablated section of a pyramidal facet is highly distinct in its smoothness. At the centre of this ablation feature was a surface roughness of $\sigma_{rms} = 3.34 \pm 0.10$Å, standard for our pyramidal mirrors and commensurate with fibre-tips. This corresponds to expected scattering losses of $\mathcal{L}_S = 29.7 \pm 1.8$ppm at 780nm, following the model of Section 3.4. While this is an order of magnitude higher than the values achieved by super-polishing, it is a very promising result considering our relatively simple method of substrate preparation. To further improve surface quality, abrasive tip polishing may be replaced in substrate manufacture. Alternatively, increased thermal annealing may be added to the fabrication process, potentially as an independent stage before ablation.

5.6.5 Dielectric Coating

To best evaluate the optical quality of the pyramidal mirrors, we instructed a mirror coating manufacturer[25] to apply a high-reflectivity dielectric coating with minimal transmission. The design wavelength was naturally $\lambda_d = 780$nm, comprising quarter-wave layers of Tantalum Pentoxide and Silicon Dioxide. These layers were applied by ion beam sputtering, with post thermal annealing to reduce contamination. The expected absorption losses of this coating were $\mathcal{L}_A < 10$ppm, which is verified in the following chapter.

To assist with this coating process, a suitable mounting fixture was developed. The schematic for this is given as Figure 5.15a). This allowed the pyramidal mirrors to be handled as a standard $\phi = 25.4$mm cylindrical optic, while containing the features

[25] Laseroptik GmbH.

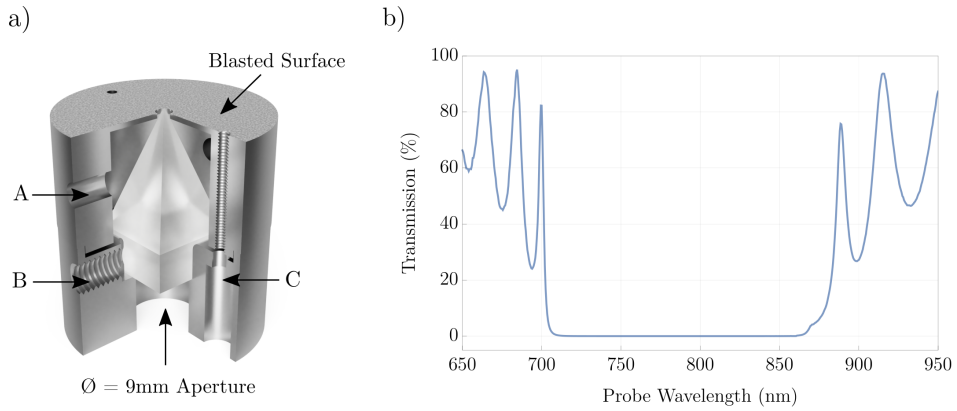

Figure 5.15: **Pyramidal Coating Fixture and Transmission. a)** A coating fixture was developed which allows for a pyramidal mirror to be handled as a standard 25.4mm cylindrical optical element. A: Ventilation holes were added for vacuum suitability. B: The pyramid was secured by spring-loaded setscrew. C: The mount could be separated into two parts, secured with screws. The pyramidal tip protruded 200μm from the sandblasted face of the mount, to allow for mirror coating. A 9mm aperture was placed at the rear of the mount, to allow for optical characterisation of the mirror surface. **b)** Coating transmission was measured by Laseroptik GmbH., using a LAMBDA 1050 spectrophotometer. This shows the broadband performance of the coating, over the entire near-infrared wavelength range. A higher resolution measurement at 780nm is performed in the next chapter.

necessary for vacuum operation and coating. Since this mount gave full optical access to the pyramidal facet, it was also used in coating characterisation. This allowed the use of standard optomechanical mounts, greatly decreasing experimental complexity.

5.6.6 Review of Manufacture

We have developed pyramidal mirrors as a potential alternative to fibre-tips, anticipating a range of advantages in cavity fabrication and operation. Glass shaping, ablation and reflective coating proved successful, with a desired range of feature geometries and mirror arrays realised. However, the use of pyramidal mirrors introduced a range of new considerations to the manufacturing process.

In shaping the glass pyramids, polishing each tip to a desired facet size provided flexibility, but may have been a limitation to final surface quality. The action of thermal

annealing inherent to ablation largely rectified this, allowing for the creation of low loss mirrors. However, the non-isotropic distribution of glass directly around each ablation site was seen to modify feature geometry. If heat was not able to dissipate into the bulk glass at a similar rate across the entire feature, local modifications to its geometry were observed. In elliptical features, if the major axis approached the edge of the facet it would experience sustained heating, increasing its eccentricity. This was reduced in using larger facet sizes, but may also be assisted by using a glass substrate that more closely matches the intended ablation feature profile. In general, this would be returning the cylindrical symmetry of fibre-tips.

To provide a comprehensive evaluation of the pyramidal substrates, their properties as a cavity mirror must be determined directly. In Chapter 6, mirror reflectivity is measured and compared to the incoherent losses predicted during fabrication. The role of micro-patterned ablation in cavity geometry is examined, verifying enhanced feature diameters and the control of cavity birefringence. In Chapter 7, a science cavity is formed using a pyramidal substrate, to act as an elementary quantum interface. A process for their alignment is developed and considered in the context of other macroscopic cavity designs. The frequency stabilisation of this cavity gives a signature of their inherent vibrational stability and susceptibility to photothermal effects.

Chapter 6

Spectral Mirror Characterisation and Multi-Feature Resonators

This chapter describes spectral characterisation of the pyramidal mirrors produced by laser ablation. Analysis of their coating quality, sources of optical loss and verification of their geometries provides key feedback to the substrate design, ablation and coating processes. This work includes the novel demonstration of multiple cavities in simultaneous resonance, formed on a single tapered substrate.

We open with the determination of pyramidal mirror reflectivity, by measurement of cavity free spectral range and linewidth. This highlights finite optical losses, which are partitioned between coating transmission and incoherent loss mechanisms. We determine the homogeneity of the applied coating and develop methods for the identification and removal of surface contamination. Mirror geometry is measured via the frequency splitting of higher order modes and the onset of clipping losses. This is compared to interferometric surface reconstruction, establishing an improved topographical fitting model for future mirror manufacture. Finally, the birefringence of the mirrors is evaluated, induced both by their elliptical geometry and dielectric coating.

6.1 Length Scanning Spectroscopy

Interferometric techniques constitute a class of the most sensitive methods for determining the optical properties of mirrors [34]. It therefore follows that mirrors are ideally

characterised by their formation into optical cavities and the subsequent observation of their resonance behaviour. It is convenient that this is the case, given our intended experimental application, as the considerations and observations outlined here form a foundation for the fabrication and implementation of science cavities considered later.

For the greatest flexibility in the range of measurements made on the mirrors, we are confined to working outside of vacuum. As such, the cavities considered will be subject to extensive vibrations transmitted through air. The fluctuations this applies to their length, when compared to the width of their resonance, are impossible to compensate for by electronic feedback. This renders many high accuracy 'two-laser' interferometric methods practically impossible [56]. We therefore employ length scanning spectroscopy, which maps the resonances of a cavity as its length is smoothly expanded and contracted. This is used to measure cavity free spectral range and linewidth, from which mirror reflectivity is determined.

6.1.1 Overview of Method

The Deschamps circle diagram of Chapter 3 outlines the fundamental geometric requirements in the formation of an optically resonating mode. However, this somewhat implies the alignment of a cavity to be the relatively simple matter of bringing two spherical mirrors sufficiently close as to produce a mode somewhere on their surface. The experimental reality for micro-mirrors is considerably more technically demanding, where the size of the mode is essentially equal to the diameter of the mirror.

The requirement to centre the cavity mode on the mirrors places a tight condition on their relative angular and transverse displacement [45]. Then, in order to excite this mode, a monochromatic laser beam must be well matched in position, size and wavefront curvature to it. Finally, for resonance to be achieved, mirror separation must be accurate to $\lambda_{probe}/2\mathcal{F}$, which for our highly reflective mirrors, is on the order of picometres.

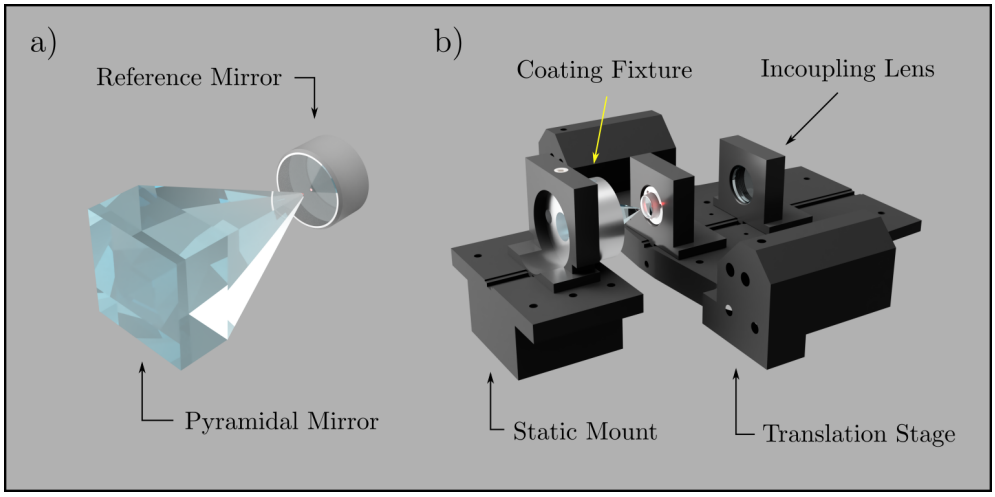

Figure 6.1: **Mirror Characterisation Cavity. a)** To determine the properties of the pyramidal mirrors, they are paired with near-planar super-polished mirrors of known optical quality. The standard cavity length range used for characterisation was 50μm to 500μm. Cavity length was set using linear micrometers, with 4mm maximum travel and a theoretical resolution of 1μm. **b)** The pyramidal mirrors are secured on a static mount, within their coating fixture. To align the cavity, the reference mirror is mounted on a six-axis translation stage. This stage includes piezoelectric actuators, to scan cavity length over $\approx 1\mu$m and create an observable resonance signal. The maximum scan range of the stage was 30μm, with a theoretical resolution of 1nm. An in-coupling lens is mounted to this stage, used to maintain illumination of the cavity mode largely irrespective of stage position.

Ultimately, the challenge of cavity characterisation comes simply as a consequence of attempting to probe an inherently optically shielded environment by optical methods. To overcome this, several established techniques have been developed and are employed here. These can be divided into two categories: those that ensure good mirror pre-alignment before the observation of resonance is attempted and those that generate an experimental signature with which it can be fine-tuned.

To ensure that a pair of mirrors are pre-aligned as to resonate when required, there are generally two schools of thought. The first is to rely upon a high machining tolerance for each component in the cavity, such that when assembled the optical axis is perfectly placed. However, this strongly limits any translational freedom that may be required,

including the use of piezoelectric elements to modulate cavity length [45]. Therefore, an approach is adopted which relies fully upon guided translations in each axis.

An illustration of the characterisation cavity is given as Figure 6.1. The pyramidal mirror of interest is paired with a near planar reference mirror of known optical quality. These reference mirrors possess high reflectivity dielectric coatings,[1] with $T = 0.5$ppm and $\mathcal{L}_T = 2$ppm, or high transmission coatings,[2] with $T = 38.2$ppm and $\mathcal{L}_T = 2$ppm. A selection is made based on preferred cavity linewidth and input optical power. The pyramidal mirror is held fixed, while the reference mirror is used to align the cavity and modify its length. The properties of the mirrors are determined by observation of resonance behaviour as length is scanned, measured in both the transmission of light and its back reflection.

6.1.2 Characterisation Apparatus and Alignment Process

The use of a planar reference mirror in cavity characterisation simplifies mode matching, as the probe beam can be focused onto it with the correct spot size, without adapting wavefront curvature. The pyramidal mirrors display the same stringent alignment tolerance as that of fibre-tips, but do not possess integrated beam steering. Accordingly, the illumination laser does not follow the translation of a pyramidal mirror, creating an additional degree of freedom in their alignment. To eliminate this from our system, the in-coupling lens was attached to the same translational platform as the input mirror.

During pre-alignment of the mirrors, optical imaging is used to ensure that the probe beam is well centred on the targeted ablation feature and back reflections are examined to ensure it runs perpendicular to each surface. Once cavity length is scanned using piezoelectric stacks integrated into the stage,[3] this is usually sufficient such that a flashing resonance appears somewhere on the feature. However, once both mirrors

[1] REO Part NS11333 Coating Run S7-249
[2] REO Part NS11334 Coating Run S7-254-P1/S7-251-P1
[3] Thorlabs MAX602D/M & Piezomechanik SVR-150

are installed, the probe beam will not be transmitted by the mirrors unless it does form a good cavity. It is therefore overlapped with a 635nm collimated laser diode, which is largely transmitted by the dielectric coating.

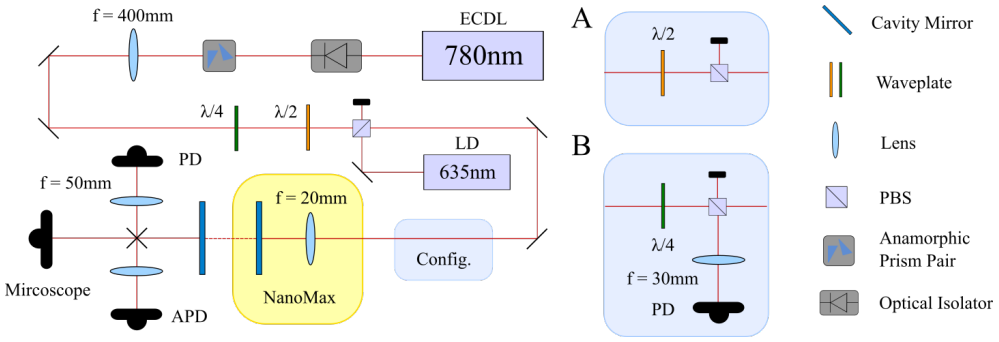

Figure 6.2: **Cavity Characterisation Beamline.** To perform length scanning spectroscopy, the cavity is illuminated with 780nm light and mirror separated is modulated using the *NanoMax* piezoelectric translation stage. To assist with alignment, a 635nm laser diode is overlapped with the probe beam. The cavity transmission signal is variably monitored on an avalanche photodiode (APD), transimpedence gain photodiode (PD) and 12x optical microscope. There are two configurations for polarisation control of the illuminating laser. Configuration A allows the linear polarisation of the beam to be rotated, for measuring polarisation mode spitting. Configuration B allows capture of the cavity back-reflection signal, for measuring mirror transmission. When not specifically stated, configuration A is used with the half wave-plate removed.

By applying curvilinear translations to the reference mirror, and fixing any observed displacement of the 635nm beam by linear corrections, the resonating mode can be moved towards the feature centre. Fine tuning of the probe beam can then be performed by examining its distribution into higher order modes. We are mostly concerned with coupling maximum power into the fundamental mode. One minor drawback of co-translating the lens with the reference mirror is that altering mirror alignment leads to an astigmatic focus. However, given the extent of corrections to alignment that are generally required in our system, the impact on fundamental mode coupling is small.

The full beamline of the system is given as Figure 6.2. For most characterisation methods it is sufficient only to examine light transmitted by the cavity. A magnetic stand

behind the pyramidal mirror allows replaceable right angle mirrors to be inserted into the beamline, directing the cavity signal to a choice of sensors. In one direction is a 2f imaging system of the cavity mode onto a silicon doped avalanche photodiode (APD),[4] for low intensity transmission measurements at high bandwidth. In the opposite route is a lower bandwidth photodiode[5] with high transimpedance gain, with a sensor size large enough to reliably address the entire cavity mode. Finally, if there is no mirror inserted, the cavity mode is imaged by a 12x variable zoom microscope.[6] This uses in-line white light illumination to resolve facet edges and therefore ablation feature position.

6.1.3 Linewidth Measurement

To determine cavity linewidth, its length is oscillated around a resonance and the transmitted intensity signal captured. This trace is fitted to a Lorentzian function and a conversion of the time axis to optical frequency is made. To make this conversion, we require known frequency markers. We achieve this by radio frequency (RF) modulation of the laser diode current, causing it to lase with sidebands on the carrier frequency. These sidebands sit at $\pm$ the drive frequency, which may be set between 0 - 120MHz using our arbitrary waveform generator.[7] The time between the cavity being resonant with the lower and higher frequency sidebands is compared to the time taken to cross resonance. An example trace is given in Figure 6.3a), where a cavity of 20MHz linewidth is addressed with 100MHz sidebands.

To ensure that the observed resonance profile is purely that of the resonator, several considerations must be made. Firstly, the spectral line shape of the probe laser is convolved with that of the cavity, requiring it to be sufficiently narrow for good resolution. The Toptica DL100 laser system used here had a linewidth of 450kHz, requiring Voight

[4] Thorlabs APD110A/M
[5] Thorlabs PDA36A-EC
[6] Navitar 1-50487 with attachments.
[7] Keysight 33600A

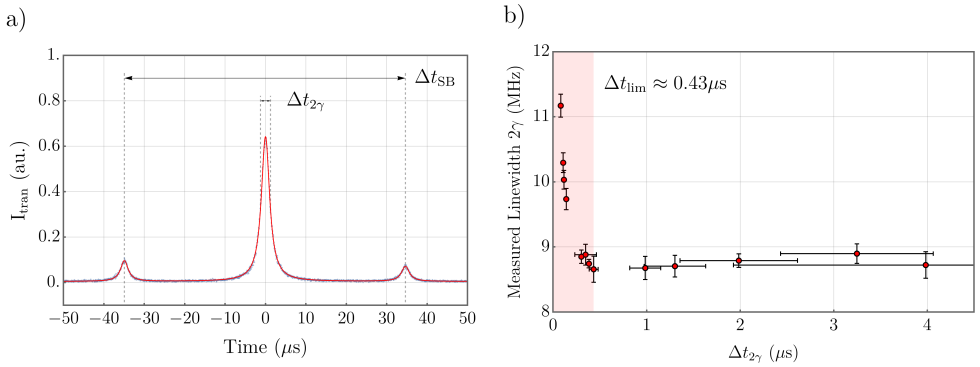

Figure 6.3: **Linewidth Measurements. a)** A representative measurement of cavity linewidth is made, showing the transmission profile of a $\Delta\nu_c \approx 20$MHz linewidth cavity illuminated by a probe beam with $\pm\Delta\nu_{SB} = \pm 100$MHz sidebands. In the time axis of changing mirror separation, these frequencies are represented by the FWHM resonance crossing time, $\Delta t_{2\gamma}$, and sideband crossing time, Δt_{SB}. **b)** The effect of detector broadening is evaluated, where resonance crossing times of $\Delta t_{2\gamma} < 0.43\mu$s are seen to inflate the measurement of cavity linewidth. This limited maximum mirror velocity, used to minimise the presence of vibrational noise in the resonance profile. This noise can be seen in the reduced error of $\Delta t_{2\gamma}$, as scan frequency is increased. While this does not visibly affect the error in measured linewidth, it is a notable limitation to the measurement of free spectral range. All errors are determined statistically.

deconvolution of the measured trace in limited cases. Then, the time taken to cross resonance must be greater than the lifetime of the cavity, as to observe only its steady-state behaviour. Finally, the crossing time must also be slower than the responsivity of the detector, which formed the most limiting aspect of our system. The relationship of mirror velocity to measured linewidth is shown in Figure 6.3b).

6.1.4 Free Spectral Range Measurement

The measurement of cavity free spectral range proceeds in a similar manner to that of linewidth, by increasing the amplitude of length oscillation to $> \lambda_{probe}/2$. An example of the observed spectra is given as Figure 6.4a). However, the sideband calibration markers of $\Delta\nu_{SB} \approx 0.1$GHz spacing are now several orders of magnitude less than the measured spectral ranges of $(1/2\pi)\Delta\omega_{fsr} = \Delta\nu_{fsr} \approx 0.5$THz. If the cavity is subject to vibrational noise with a frequency spectrum that alters sideband crossing time without

a proportional change to spectral range crossing time, significant error is introduced into the measurement results. Indeed, such vibrational noise has already been observed in Figure 6.3b), although for the measurement of linewidth the sideband crossing time was sufficiently close to the linewidth crossing time as to be equally affected by noise.

In contrast, for the measurement of free spectral range, mirror velocity must be carefully chosen to minimise calibration error. Notable mechanical resonances also exist within the stage, rendering velocity a complex function of driving voltage. A demonstration of this behaviour and its optimisation are given as Figure 6.4b). It is important to further note that since the flexure stage exhibits significant cross-talk between the axes of translation, any shift in mirror position alters this profile.

A further issue in the accuracy of spectral range measurements is that the piezo extension is non-linearly proportional to its applied voltage. Therefore, the linear proportionality assumed between crossing time and frequency separation in linewidth determination must now be rectified. The solution was to map frequency scanning rates around each transverse mode, building coordinates of $(t, \Delta\nu_{SB}/\Delta t_{SB})$. A 4^{th} order polynomial fit was made to these coordinates, of which an example is given as Figure 6.4c). The free spectral range is then simply an integration under this curve, between the times of fundamental mode resonance.

The integral approach did not yield accurate results in all cases. Occasionally, the nonlinearity profile would be such that a polynomial fit was not appropriate. Further, cavities near the point of confocality exhibit an insufficient spacing of transverse frequencies to map nonlinearity. Finally, the characterisation of cavities subject to clipping loss displayed too few low-loss higher order modes to apply a fit. These issues were overcome by using free spectral range measurements to calibrate the high resolution micrometer of the translation stage to cavity length. The accuracy of this approach is shown in Figure 6.4d), where outlying measurements of cavity length can be easily

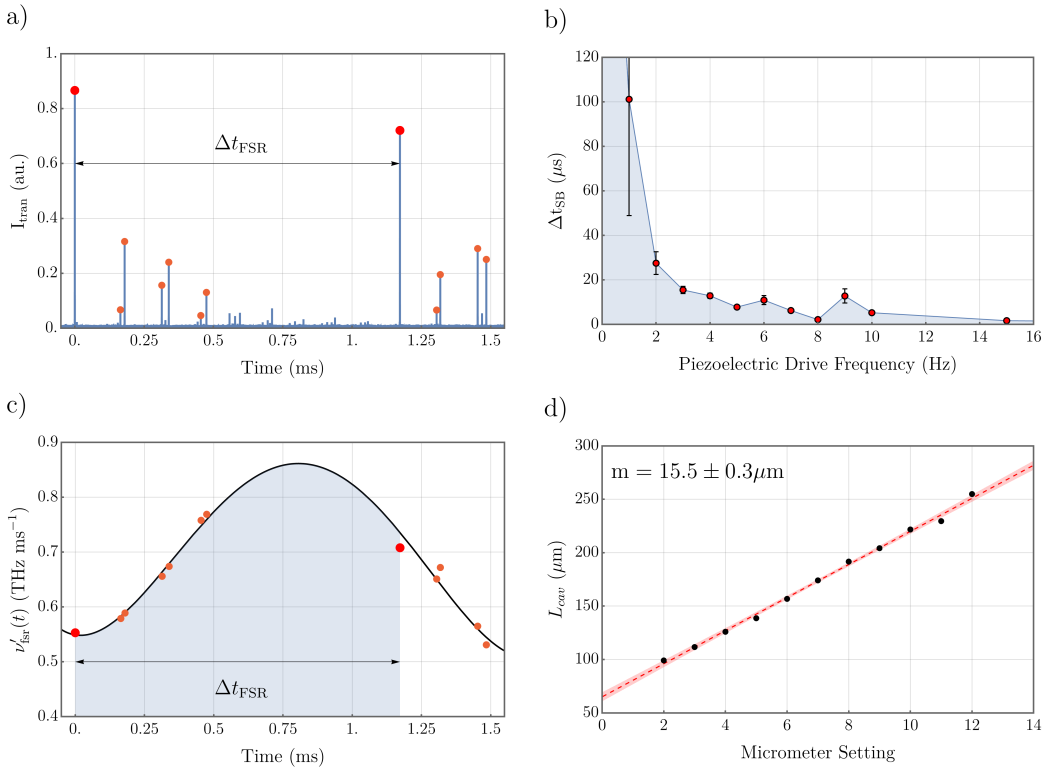

Figure 6.4: **Free Spectral Range Measurement. a)** $\Delta\nu_{fsr}$ measurement; timing consecutive fundamental resonances (red), with higher order modes (orange). **b)** Due to vibration, slow sideband crossings were subject to error, reducing calibration accuracy. At higher frequencies, detector broadening reduced sideband resolution. The ideal characterisation frequency was at around 8Hz. **c)** To compensate for piezo nonlinearity, the frequency scanning rate was fitted for every mode, mapping mirror velocity. Spectral separations are the integral region, where $\Delta\nu_{fsr}$ is highlighted. **d)** Since higher-order modes were not always resolvable, the high resolution micrometer was calibrated to cavity length using available $\Delta\nu_{fsr}$ measurements.

identified and corrected.

6.1.5 Review of Technique

The use of a planar reference mirror and co-translating lens largely mitigated the challenge of pyramid alignment, when compared to the characterisation of a paired fibretip cavity. When combined with off-resonant illumination of the mirrors and the ability to observe the cavity mode profile via the pyramid, the formation and optimisation of a cavity resonance signal proved time efficient and highly repeatable.

The chosen method of linewidth determination, by scanning cavity length around resonance with a sideband modulated laser, is common within the field. Our standard error for this measurement is $\approx$ 1-4%, determined by statistical uncertainty. However, our process of free spectral range measurement is far less common, with more favoured approaches using two-laser spectroscopic methods [122] or imaging cavity length with a calibrated microscope [47]. Dual laser illumination is noted for its higher precision, but requires the suppression of vibrational noise on the cavity. The use of an optical microscope is far less experimentally complicated, but was prevented by our pyramidal facet size, $< 350\mu$m, being significantly smaller than the planar mirror diameter, 7.75mm. A standard microscope was therefore unable to focus on the edge of each mirror simultaneously, to accurately determine their separation.

The drive frequency of cavity length scanning was carefully tailored to reduce the impact of environmental noise. Piezo nonlinearity was corrected by the mapping of higher order mode resonances and calibration of a high resolution micrometer. This created a standard statistical error of $\approx$ 2-5% for a given measurement, which was perfectly sufficient for the characterisation work to follow. Since our method of determining linewidth and free spectral range are essentially identical, modifying only the cavity length scan, the measurement of mirror reflectivity proved a rapid process.

6.2 Dielectric Coating Evaluation

In this section we perform spectral analysis on the dielectric coating applied to the pyramidal mirrors. This will allow independent evaluation of factors arising from both the ablation and commercial coating processes. We will compare the spectrally measured properties to those predicted by surface analysis and coating specifications; with any discrepancies forming the basis of feedback to future manufacturing cycles. We shall chiefly consider three examples of pyramidal mirrors, representing a range of ge-

ometries. For consistency, these are the mirrors given in Figure 5.13, with the same labelling used here.

	A		**B**		**C**
R_{oc}	800μm		1200μm		250μm
Ref.	HR (V)	HR (H)	HR (H)	OC (H)	OC (H)
Finesse	158,817	143,744	151,323	79,207	81,575
$\mathcal{L}+T$ (ppm)	39.56	43.71	41.52	79.33	77.02
$\mathcal{L}_{Pyr}+T_{Pyr}$ (ppm)	39.14		39.02	39.13	38.82

Table 6.1: **Pyramidal Mirror Losses.** We consider the reflectivity of three pyramidal mirrors, whose surface profiles were previously given as Figure 5.13. For pyramidal mirror C, the high curvature feature in the array is discussed. Two types of reference mirrors were used: OC with $T = 38.2$ppm, $\mathcal{L}_T = 2$ppm and HR with $T = 0.5$ppm, $\mathcal{L}_T = 2$ppm. These properties were independently verified by ring-down spectroscopy. Two polarisations were used, horizontal (H) and vertical (V) with respect to the optical table. $\mathcal{L}+T$ are total losses of the cavity and $\mathcal{L}_{Pyr}+T_{Pyr}$ are the remaining losses when those of the reference mirror are subtracted. A mean value is given for pyramidal mirror A, from its measurement both in (H) and (V).

6.2.1 Reflectivity and Total Losses

Mirror reflectivity was determined via cavity finesse, taking the ratio of cavity free spectral range to its linewidth. Measurements were taken for a range of mirror geometries and reference mirrors, with selected results displayed in Table 6.1. Excellent consistency was seen in these results, implying pyramidal mirror losses of approximately $\mathcal{L}_T + T = 39$ppm. In the previous chapter, we anticipated scatting losses of $\mathcal{L}_S \approx 29$ppm and absorption losses of $\mathcal{L}_A < 10$ppm. As we are at the upper end of our prediction, we investigated the source of these losses further. We are chiefly interested in the contribution of transmission to measured losses, as well as the potential presence of surface contamination or coating inhomogeneity.

6.2.2 Transmission

Transmission of light through the mirrors reduces cavity finesse in the same manner as scattering and absorption losses. An accurate measurement of transmission will therefore quantify the extent of incoherent losses in the pyramidal mirrors, as well as highlight the current state of the art in commercial dielectric coating application. For this batch of mirrors the transmission of the dielectric coating was designed to be minimal, as to benchmark the best reflectivity available with our substrates. However, some finite transmission is always expected to be present in dielectric coatings, as a result of variation in stack layer thickness. Even residual transmission is crucial for optical interrogation of the cavity, without which the work of this chapter would be impossible. Spectrophotometry performed at the point of manufacture[8] indicated a residual transmission of $T < 15$ppm, but given the highly attenuating nature of the media, more precise analysis was not practical via this method. Instead, this was achieved using a standard cavity-based analysis [35], comparing transmitted and back-reflected cavity signals.

To understand this measurement we return to Equation 3.8, describing the intensity of light transmitted by the cavity on resonance. This allows mirror transmission to be calculated directly from total cavity losses and reference mirror properties. However, this expression neglected to consider the geometrical overlap of the probe beam with the cavity mode. In any practical scenario, a notable portion of incident light does not couple into the mode and accordingly reduces transmitted light. To rectify this, we introduce a mode-matching factor, ϵ, such that the power of incident light coupled into the cavity is given by ϵP_{inc}. Our modified transmission is then:

$$\left.\frac{P_{tran}}{\epsilon P_{inc}}\right|_{\omega_{res}} = \frac{4T_1T_2}{(T_1+T_2+\mathcal{L}_1+\mathcal{L}_2)^2}. \tag{6.1}$$

[8] Using a PerkinElmer LAMBDA 1050 spectrophotometer.

Of the laser power incident on the cavity, $(1-\epsilon)P_{inc}$ will be immediately retroflected without the possibility to couple into the cavity mode. This will not interfere with light exiting the cavity and thus decrease the visibility of the back-reflected 'dip' on resonance. The power ϵP_{inc} will continue to destructively interfere with exiting light, whose cancellation is given by the impedance matching of the system. Overall, this gives a modified expression of:

$$\left.\frac{P_{ref}-(1-\epsilon)P_{inc}}{\epsilon P_{inc}}\right|_{\omega_{res}} = \left(\frac{T_2+\mathcal{L}_1+\mathcal{L}_2-T_1}{T_1+T_2+\mathcal{L}_1+\mathcal{L}_2}\right)^2. \tag{6.2}$$

By subtraction of one from Equation 6.2 and division of Equation 6.1 by the result, we derive an expression that does not involve geometrical mode-matching. Thus, from the simultaneous measurement of the power incident, transmitted and reflected by the cavity, pyramidal transmission can be determined. This was performed for Pyramidal Mirror B, using an OC reference mirror.

For measurements of output power, large area photodiodes were used to ensure full beam capture. The transmitted beam was routed towards this sensor using the magnetic switching described in Section 6.1.2. Consecutive measurement of cavity linewidth using the APD and PD verified that the reduction in sensor bandwidth did not affect results. The input and back-reflected light of the cavity were separated by placing a polarising beam-splitter and quarter wave-plate before the cavity. Imperfect mode-matching necessitated high illumination power, so a variable ND filter was introduced to control sensor exposure. The beamline for these measurements is shown as configuration B in Figure 6.2. To measure cavity input power, account for optical losses within the system and ensure the calibration of each photodiode, power was sampled between each optical element by a handheld meter.[9]

Resonant intensity maxima and minima were determined directly by the oscilloscope,[10]

[9] Thorlabs PM100A & S130C
[10] Keysight MSO9254A

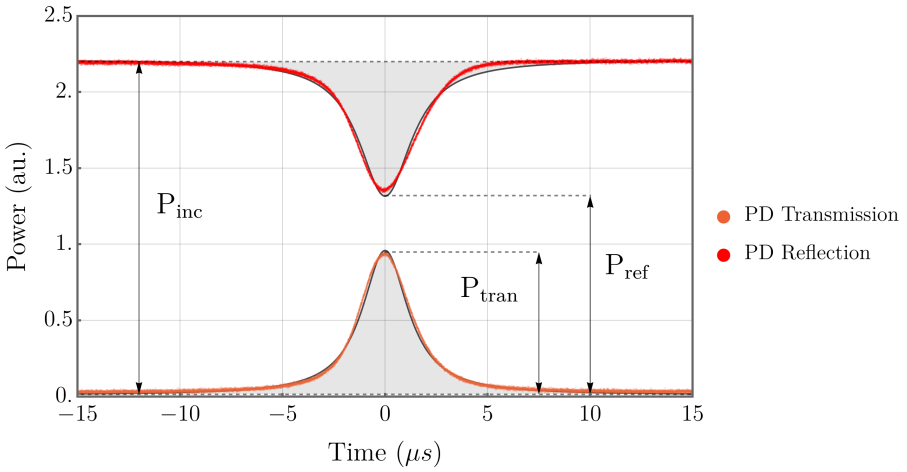

Figure 6.5: **Transmission Measurement.** To derive geometric mode matching, the optical power transmitted by the cavity is compared to the cancellation of back reflected light. This allows a determination of mirror transmission. An example measurement is shown in colour, with fitted Lorentzian profiles in black. Note that these are raw photodiode captures to allow similar scaling, with true power values obtained by factoring in systematic losses.

by fitting 500 Lorentzian line-shapes to consecutive cavity length extensions. The statistical error proved to be greater than the fitting error, caused by power fluctuations of the probe laser and optical interference from the facets of the beam-splitter. An example of a measured trace is given as Figure 6.5. The input power was measured to be $P_{in} = 0.36$mW. This gave output powers of $P_{tran} = 5.25\mu$W and $P_{ref} = 237.84\mu$W, corresponding to a mode matching efficiency of $\epsilon = 0.34$. Of the 39ppm total losses attributed to the pyramidal mirror, 1.74 ± 0.23ppm were therefore determined to be in transmission. While this is greater than our $T = 0.5$ppm HR super-polished reference mirrors, it remains significantly smaller than incoherent loss mechanisms.

An alternative method of calculating cavity mode-matching is to make use of the observed transverse mode spectrum as a decomposition of the spatial profile of the probe beam into that of the cavity modes. Mode-matching is then the ratio of optical power coupled into the fundamental mode, compared to all higher orders. However, this measurement can suffer from multiple issues of accuracy. Higher order modes have

corresponding larger mode sizes, so experience greater loss on small diameter mirrors. Then, any inhomogeneity in the mirror coating or scattered impurities would variably affect their finesse. To avoid these concerns, mode-matching was only determined using the first method outlined.

6.2.3 Coating Homogeneity

The homogeneity of the dielectric coating on each pyramidal mirror is a reliable indicator of a range of manufacturing errors, as well as a potentially limiting factor in the fabrication of science cavities. It can be examined by observing the relationship between cavity mode position on the mirror of interest and its finesse. If a reduction in finesse is associated with a gain in transmission, then the curvature of the feature may be acting to warp the thickness of the dielectric layers. If losses are isotropic and purely incoherent, then the ablation parameters may not be smoothing the glass surface over a sufficient area. Finally, if losses are highly localised, contamination or damage of the mirrors may be present. This is known to give a distribution of finesse values between higher order transverse modes, as considered in the following section.

To evaluate the homogeneity of our mirror coating, the cavity mode was walked across the pyramidal mirror facet and reflectivity measured at each position. For the semi-planar cavity geometry employed, mode migration is simply a case of changing the pitch and yaw of the reference mirror. The length of the cavity must be carefully chosen by considering the maximum range of translation stage. A shorter cavity leads to a smaller mode size on the pyramidal mirror and thus a higher resolution sampling of the surface. However, a longer cavity has a greater mode migration for the same change to pitch or yaw, controlling the extent of the surface that can be sampled. For our presented measurement of Pyramidal Mirror B a cavity length of 300μm was selected, sufficient to sample a 100μm$\times$100μm area with a spot size of 14μm.

The position of the fundamental mode on the surface was tracked using the calibrated

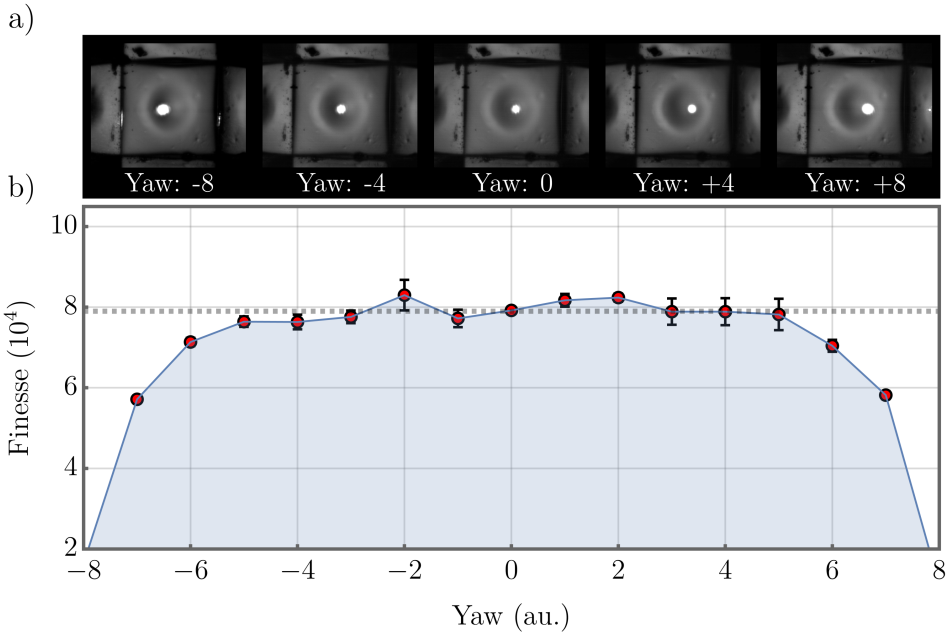

Figure 6.6: **Coating Homogeneity. a)** Cavity mode position was migrated horizontally across the surface of Pyramidal Mirror B, monitoring mirror reflectivity. Images of fundamental mode resonance were captured for different yaw positions of the planar reference mirror. **b)** The changing reflectivity of the 100μm surface is indicated in its finesse. As the mode approaches the feature edge clipping losses are visible, however excellent homogeneity is seen at the mirror centre. The waist of the mode addressing the mirror was 15μm, acting as a Gaussian filter to the results.

zoom microscope and migrated across the full extent of the mirror in its vertical and horizontal axes. A demonstration of the observed change in mode position is given as Figure 6.6a). The relative simplicity of this measurement is valuable aspect of the use of pyramidal mirrors. Such a technique could not be applied to a fibre-tip micro-mirror as imaging the position of the cavity mode on its surface would be considerably more challenging. The linewidth of the cavity was measured at each position to quantify optical loss. Any change to the FSR induced by translational cross-talk was monitored both spectrally and by using a secondary microscope.

The homogeneity of the coating in the horizontal axis is given as Figure 6.6b). This has the profile that would be expected from a good quality mirror, where there is only a

small reduction to cavity finesse towards the boundaries of the feature. Indeed, the only notable source of losses is the rapid onset of clipping as the mode is translated over the feature edge. Overall, this gives us greater flexibility in cavity fabrication, as mirror reflectivity is less sensitive to cavity mode size and more tolerant to misalignment. However, in the vertical axis the higher order modes demonstrated lower linewidth, as shown in Figure 6.8. As this is a known indication of mirror contamination, a detailed investigation is performed in the following section.

6.2.4 Identifying Contamination

In addition to vibrational instability, a further issue of characterising mirrors outside of vacuum is the presence of airborne particles. A single particle of contamination would occasionally adhere to an exposed dielectric surface and induce highly localised optical losses. In order to distinguish losses induced in this way, to those related to manufacture, a range of measurements were developed. Then, a verified method for their removal was established.

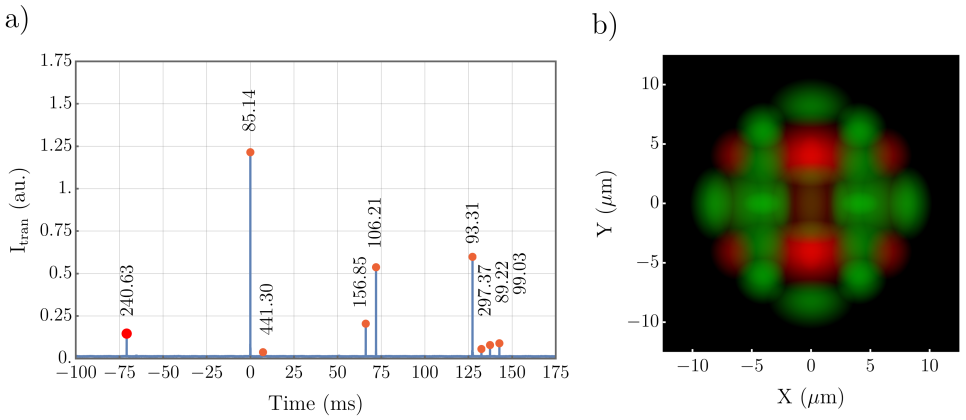

Figure 6.7: **Static Loss Analysis with Higher Order Modes. a)** Transverse mode spectrum for a mirror with particulate contamination. Each spatial cavity mode is shown with its associated total losses in ppm. **b)** Loss visualisation tool, highlighting the spatial distribution of low loss modes in green and high loss modes in red. Losses are seen to be highly localised and likely caused by an off-centre speck of dust, located along the vertical axis.

The most readily apparent signature of particulate contamination was a difficulty in maximising laser coupling into the fundamental cavity mode, as measured by its transmission. For a reasonably aligned probe beam, one would expect a smooth decrease in peak transmission as transverse mode number is increased. However, the transverse mode spectrum of a cavity subject to particulate loss displayed a notable suppression of the fundamental mode. This was explained by each higher order mode cluster containing at least one solution with a nodal line along the area of contamination. This mode accordingly experienced lower optical losses and transmitted greater power. In contrast, the fundamental mode has no degeneracy or nodal lines with which to avoid contamination. An example of such a spectrum is shown in Figure 6.7a).

To more clearly distinguish between the impacts of mode-matching and particulate losses on the transverse mode spectrum, the linewidth of each observed mode was fitted. Indeed, the fundamental mode was seen to have a broadened linewidth, with a complex distribution of losses amongst higher orders. As a visualisation tool, each mode profile was coloured according to its derived losses and overlapped. This yielded a probability map given as Figure 6.7b), which indicated the potential locations of high loss. This method shares interesting parallels with an established technique of intracavity atomic position measurement [137].

There were several issues with this approach, highlighted by its use to locate a single piece of dust on a cavity mirror. The first was that modes subject to high losses would simply not be present in the acquired spectrum, making accurate normalisation of the probability map impossible. Further, given the highly structured and spatially symmetric profiles of higher order modes, probabilities were usually distributed over a number of distinct lobes.[11] This was partially improved by the ellipticity of the pyramidal mirrors, which split the degeneracy of higher order mode clusters and allowed for their independent characterisation. However, the location of the single particle in this

[11] This leads to the informal description of this method as a 'cavity kaleidoscope'.

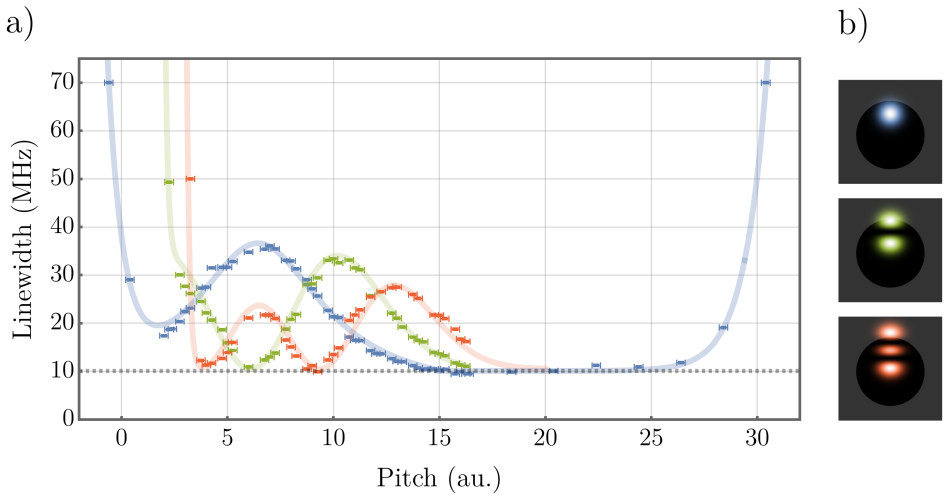

Figure 6.8: **Localising Contamination. a)** Using the technique of Section 6.2.3, cavity mode position was migrated across the vertical axis of Pyramidal Mirror B, recording the linewidth of three transverse modes. The change in linewidth of these modes are shown, where the dust particle is located at the maxima of the blue trace. The expected linewidth of the cavity, from characterisation of clean pyramidal mirrors, is shown as the grey line. **b)** The transverse profile of each mode is shown, with the two higher orders having nodal lines of zero intensity. As can be seen from a), when the dust sits on one of these nodal lines the minimum linewidth is achieved. This allows many higher order modes to effectively ignore contamination, resulting in their higher finesse.

case could only be determined to a position either directly above or below the mode centre.

A superior characterisation is an expansion of work of Section 6.2.3, by walking the cavity mode across the pyramidal mirror, now resolving the linewidth of each transverse mode. This constitutes a novel optical method for generating a two-dimensional loss map of a mirror surface. The results for the vertical axis of Pyramidal Mirror B are given as Figure 6.8. The position of the particle is now clearly indicated by the peak fundamental mode linewidth. Interestingly, since the modes are considerably larger than the contaminant, this measurement has essentially imaged their intensity distributions. When the particle of dust was located along a nodal line of a higher order transverse mode profile, the reflectivity expected of an uncontaminated coating

was measured. This highlights the importance of considering transverse mode order when performing more basic characterisations of mirror reflectivity, as losses can be easily overlooked.

This mirror was later observed under a 24x optical microscope and a particle of dust was seen at the expected position. To remove this contamination, First Contact polymer was used. This was preferable to solvent cleaning, as there was no requirement to develop drying methods that prevented streaking. This was a notable advantage of the pyramidal substrates over optical fibres, whose fragility prevents cleaning in this manner. Following this process, the expected cavity finesse was re-achieved.

6.2.5 Review of Coating

The measurements of this section confirmed the successful application of a highly reflective dielectric coating to the pyramidal mirrors. The losses of this coating were consistently measured as $\mathcal{L}_T + T = 39$ppm, corresponding to a reflectivity of $R \approx 99.996\%$. Of these losses, approximately 2ppm were determined to be in transmission. This is consistent with the surface scattering losses predicted by atomic force microscopy, $\mathcal{L}_S = 29$ppm, when the residual losses, $\mathcal{L}_A = 8$ppm, are assumed to be in coating absorption. This corresponds well with the state of the art in laser ablation, with external studies reporting incoherent losses between 26ppm and 40ppm [41, 44, 49, 59, 138].

The most effectual improvement to mirror quality would be to reduce the roughness of the pyramidal facets before ablation. These were prepared using 20nm grit paper, smoothed to $\sigma_{rms} \approx 3$Å optical quality by the process of thermal annealing. An improved approach may be to start with a planar super-polished substrate and grind its edges to a taper. However, this gives less flexibility in altering the dimensions of the facet, requiring a more definite cavity design at the time of substrate production.

Alternatively, the pulse sequence used in ablation may be further optimised to promote

surface healing. It is notable that no correlation between feature geometry and surface quality was observed in our measurements, nor any inhomogeneity in mirror reflectivity across their surface. This is unusual, as a relationship between feature depth and surface roughness was expected, with a more significant molten layer corresponding to lower optical losses. It therefore seems a prudent strategy to evaluate thermal annealing as an independent fabrication stage in future work, performed before ablation.

6.3 Feature Geometry

Throughout the laser ablation process, the evolving geometry of each feature was monitored by white light interferometry. This proved essential in fabricating mirror geometries of choice, via the iterative tailoring of pulse sequences. However, outstanding questions remained on the extent to which the Gaussian depressions created could be considered as a spherical mirror. In this section we examine if the geometry indicated by interferometric reconstruction is representative of the optical behaviour of each pyramidal mirror. The geometries of a range of samples are determined spectrally and compared to interferometry performed at the point of manufacture.

Mirror curvature is measured from the frequency splitting of higher order modes and the onset of clipping losses as cavity length is increased. Clipping losses are additionally used to determine the effective diameter of our mirrors and provide a direct indication of maximum pyramidal cavity length. Cavity birefringence is quantified by polarisation mode splitting of the fundamental mode and compared to mirror ellipticity.

6.3.1 Curvature

Feature curvature was measured via the frequency splitting between fundamental, $HG_{0,0}$ and first order, $HG_{m+n=1}$, transverse modes. This was performed using the same method as for measuring free spectral range, as described in Section 6.1.4. Indeed, an identical cavity length scan was captured for these measurements, to allow

for the compensation of piezo nonlinearity. However, the final integration range was naturally reduced, given the smaller frequency spacing of transverse modes with respect to longitudinal modes. By manipulation of Equation 3.14, the splitting between these orders is given by:

$$\Delta\omega_t = \omega_{res,n+m=1} - \omega_{res,n+m=0} = \frac{c}{L_{cav}} \arccos\left(\sqrt{g_1 g_2}\right). \tag{6.3}$$

For a cavity of known length and reference mirror curvature, pyramidal mirror curvature can be determined as:

$$R_{oc,1}(L_{cav}, \Delta\omega_t) = \frac{L_{cav}(L_{cav} - R_{oc,2})}{L_{cav} - R_{oc,2} + R_{oc,2}\cos^2\left(\frac{\Delta\omega_t L_{cav}}{c}\right)}. \tag{6.4}$$

All tested features exhibited some degree of ellipticity, splitting the frequencies of their HG_{01} and HG_{10} modes. Their measurement therefore gave two curvatures, corresponding to the major and minor axes of the mirror. Figure 6.9 shows the measurement of curvature for Pyramidal Mirror B, determined to be $R_{\mathrm{oc,V}} = 1.42 \pm 0.06$mm and $R_{\mathrm{oc,H}} = 1.70 \pm 0.09$mm in its vertical (V) and horizontal (H) axes respectively. This corresponds to an eccentricity of $\varepsilon = 0.41$, where:

$$\varepsilon = \sqrt{1 - \frac{R_{\mathrm{oc,V}}}{R_{\mathrm{oc,H}}}}. \tag{6.5}$$

The measurement of transverse mode splitting was subject to a similar source of calibration error as those of free spectral range. That is, the frequency difference between transverse modes, nominally 100GHz, is considerably larger than the sideband calibration spacing of 200MHz. The corresponding difference in scanning transit time therefore renders them non-proportionally modified by an applied vibrational noise spectrum. The measurement of transverse mode splitting was subject to notable error, as it closely corresponded to a peak of the scanning stage noise spectrum. We therefore measured splitting over a range of cavity length values, taking a mean average as our final result.

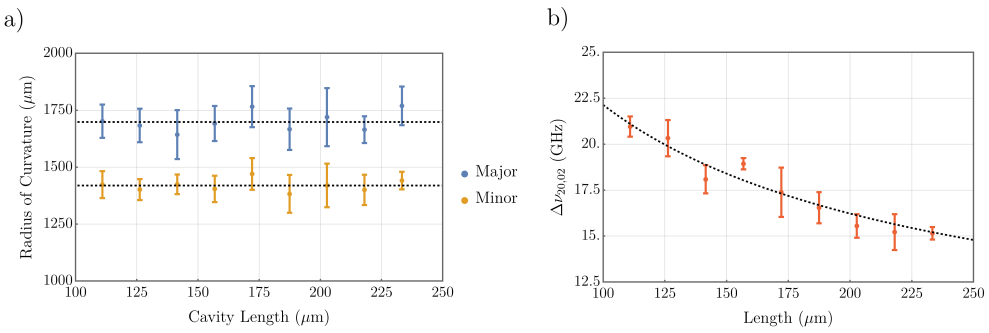

Figure 6.9: **Curvature Measurement. a)** Measurement of curvature of Pyramidal Mirror B by first order transverse mode splitting from the fundamental. The mean of each axial curvature is shown as a dashed grey line. Measurement uncertainties are determined statistically. **b)** Splitting between first order modes, as a function of cavity length. The relationship predicted by the mean values of a) are shown here as a dashed line, showing consistency of the measured results.

To further verify these results, we measured the splitting between the HG_{01} and HG_{10} modes at each length value. This splitting had a nominal value of 11GHz and was accordingly subject to significantly less noise. An example of these measurements is given as Figure 6.9b), plotted alongside the theoretical relationship expected from the results of Figure 6.9a). As can be seen, the two measurements are highly consistent with each other.

6.3.2 Diameter and Clipping Losses

A standard method for measuring the effective diameter of an ablation mirror is by observing the onset of clipping losses when increasing cavity length [44, 47]. This loss mechanism was discussed in Section 4.4, in the context of outlining the maximum practical cavity length for coherent atomic interaction. To review, as cavity length is expanded the size of the cavity mode on the pyramidal mirror grows, eventually surpassing its effective area. This rapidly degrades finesse by clipping loss, with a decay profile representative of mirror diameter and curvature.

In order to distinguish this source of incoherent loss from those considered previously, the finesse of multiple transverse modes were monitored. Since higher order modes have progressively larger areas than the fundamental, they experience the onset of clipping losses at shorter cavity lengths. Therefore, the successive decay of each transverse mode was noted as a key signature of clipping. An example of this behaviour is given as Figure 6.10a). This shows the onset of clipping for four spatial modes, measured using a single-shot ablation feature. The inherently restricted diameters of single-shot features, $\sigma_m \approx 40\mu$m, led to the onset of clipping losses at the shortest cavity lengths observed.

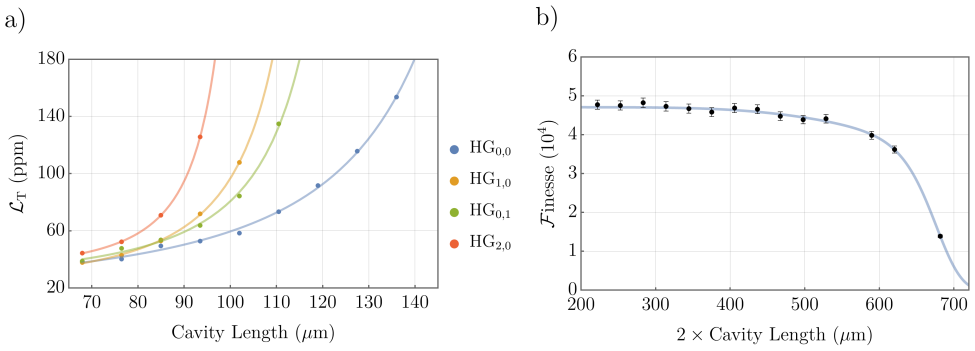

Figure 6.10: **Clipping Loss Onset.** **a)** The onset of clipping losses are observed for a selection of transverse modes. A single-shot ablation feature was used, with geometry of $\sigma_m \approx 37\mu$m and $R_{oc} \approx 250\mu$m. As can be seen, a key signature of clipping loss is its staggered onset towards progressively lower HG_{n+m} order transverse modes. **b)** The onset of clipping losses for the fundamental mode of Pyramidal Mirror A is shown. Its geometry can be derived from fitting a theoretical curve to the data, shown as a blue line. Using a dual-feature model, this gave values of $\sigma_{m,x} = 51.65\mu$m, $\sigma_{m,y} = 62.34\mu$m, $R_{oc,x} = 837\mu$m and $R_{oc,y} = 738.85\mu$m.

Pyramidal mirrors were characterised in the standard semi-planar geometry using reference mirrors. Since the diameter of these reference mirrors were considerably larger than the facet size of the pyramids, $\sigma_{m,ref} \approx 7.75$mm, all clipping losses were be attributed to the pyramids. However, trial fitting of the measured data was found to be most effective using a dual curvature clipping loss model:

$$\mathcal{F} = \frac{2\pi}{\mathcal{L}_A + \mathcal{L}_S + \frac{1}{2}\sum_i \mathcal{L}_{\mathrm{C}}(L_{cav}, R_{\mathrm{oc},i}, \sigma_{m,i})}, \tag{6.6}$$

where clipping losses arise from mirror geometries labelled $i = \{1, 2\}$. The rationale supporting the use of this model is the elliptical mirrors having a major and minor axis of curvature. As a demonstration, the onset of clipping losses for Pyramidal Mirror A is fitted in Figure 6.10b). Here, the cavity length values measured via its free-spectral-range are doubled, to emulate a geometrically symmetric mode.

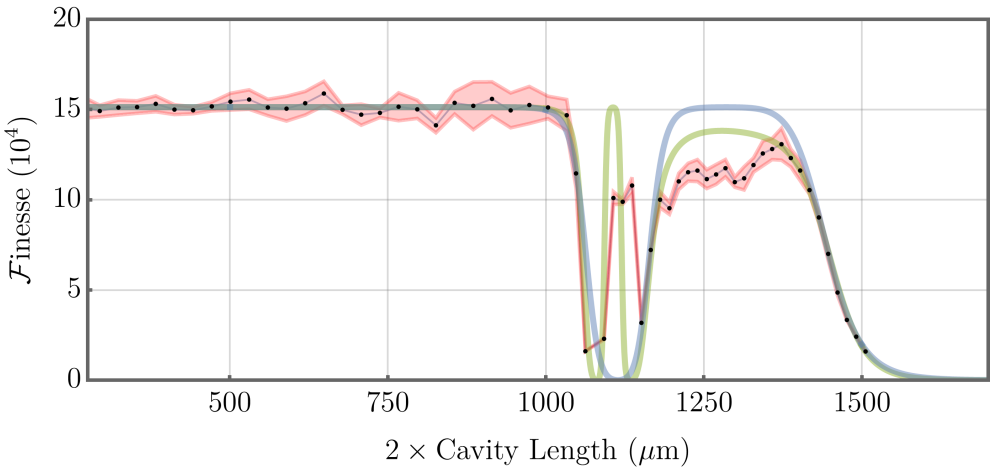

Figure 6.11: **Micro-Patterned Feature Diameter.** The onset of clipping losses was measured for a full micro-patterned ablation feature. This had no detectable reduction to cavity finesse for lengths less than 1mm, corresponding to an effective mirror diameter of $\sigma_m > 98\mu$m. For length values greater than this, the cavity exhibited complex resonator instabilities. Shown as a blue trace is a two-geometry fit, with parameters of $\sigma_{m,x} = 120\mu$m, $\sigma_{m,y} = 93\mu$m, $R_{oc,x} = 1114\mu$m and $R_{oc,y} = 1740\mu$m. To better reconstruct the finesse revival in the first instability region, a three-geometry fit was applied, with parameters of $\sigma_m = 98\mu$m, $R_{oc,1} = 1080\mu$m, $R_{oc,2} = 1134\mu$m and $R_{oc,3} = 1702\mu$m, shown in green. However, this is not a physical approximation.

The most notable result was the characterisation of Pyramidal Mirror B. The ablation sequence for this feature was intended to maximise mirror diameter, by using the micro-patterning method discussed in Section 5.3.4. The results of this characterisation are

given as Figure 6.11, highlighting a 1mm long cavity[12] with no degradation to finesse. This corresponds to an effective mirror diameter of $\sigma_m \approx 100\mu$m, reaching the current state of the art [47]. This extended cavity will allow for the integration of an optical dipole trap to our cavity system, as discussed in the following chapter.

The degradation of finesse in this cavity presented as resonator instability, given by its rapid decay and subsequent revival. Further, regions of low finesse in the fundamental mode were not associated with lower finesse in higher orders. However, the profile of this revival was complex, with features not fully described by the two geometry loss model. A three-loss model was trailed, which successfully reproduced the revival feature observed at $L_{\text{cav}} = 1150\mu$m. However, three geometrical mirror components is highly non-physical, suggesting a more obscure mechanism at work.

It is likely that the Gaussian profile of the mirror is responsible for its complex clipping loss profile. This is known to generate transverse mode coupling in microcavities, increasing the diffraction losses experienced by the fundamental [42, 43]. However, since our system achieved the requisite length for dipole trapping this was not investigated further. A second interesting extension to this work would be to simulate cavity properties from directly measured surface profiles, as has been performed for similar structures using numerical Fourier transform methods [47, 139].

6.3.3 Comparison to Surface Topography Measurements

We are now in a position to make a direct comparison between the spectrally measured geometries of the pyramidal mirrors and those predicted by white light interferometry. Ultimately, consistency was achieved between surface maps produced by phase scanning interferometry at the time of ablation and the measurements performed in this section. However, for a given pyramidal mirror, different values of curvature were indicated by

[12] This quoted length is given by the effective mode geometry for a symmetric cavity. However, since a planar mirror was used, the true cavity length was $L_{cav} = 0.5$mm.

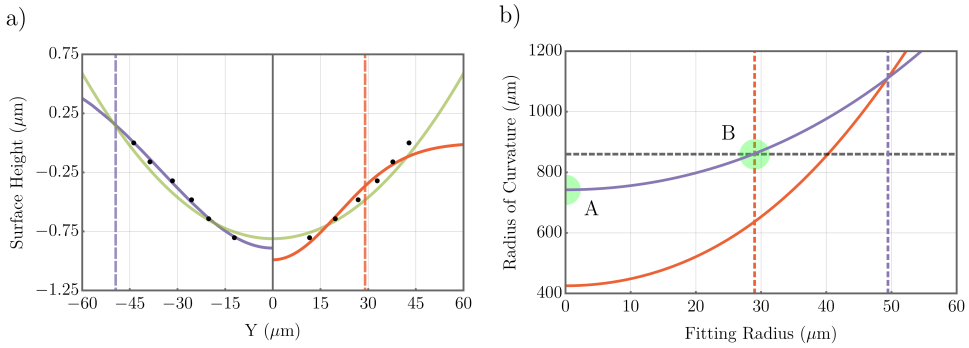

Figure 6.12: **Interferometric Reconstruction. a)** From reconstruction of the interferometric pattern given as Figure 5.13a), topographical contours on the surface of Pyramidal Mirror A are plotted. These are fitted to a Gaussian depression, either constrained to flatten at the highest contour (orange) or unconstrained (purple). Their respective 1/e radii are given as vertical lines. **b)** Mirror curvature is estimated by applying a spherical fit to the central portion of these Gaussian profiles, extending to a given radius. A: Clipping loss measurements are in agreement with unconstrained curvature values at the feature centre. B: Higher order mode splitting measurements (horizontal line) are also in agreement with unconstrained curvature values, now taken at the 1/e radius of the constrained fit.

clipping loss onset and higher order mode splitting. For a single surface map to account for each, the process of fitting a spherical profile to the mirror depression had to be treated with care.

A common issue with the accuracy of interferometric reconstruction is deciding the area over which a near-Gaussian ablation feature can be approximated as 'spherical', for the purpose of applying a fit. There is no clear consensus in academic literature for this, with studies variably taking the curvature at the centre of the ablation feature [40, 44] or fitting to its 1/e Gaussian diameter [47, 124]. As the diameter of the Gaussian profile used to perform the spherical fit is expanded, the indicated radius of curvature of the mirror increases. An example of this relationship is given as Figure 6.12b), using Pyramidal Mirror A. Here, a Gaussian depression profile was fitted to phase scanning interference extrema, of which a central portion was used to generate a secondary spherical fit.

An issue experienced with the pyramidal mirrors is that they generally possessed soft curvatures, whose boundary is occluded by curvature on the substrate edges. As a result, interferometric reconstruction does not yield a well defined surface position into which the Gaussian depression is made. Problematically, an accurate determination of the depth of a Gaussian feature is crucial in estimating its 1/e diameter. To overcome this, two independent fits are performed on a measured interference pattern. First is an unconstrained Gaussian fit, which matches the curvature of the feature well, but generally gives a non-physical estimation of its 1/e radius. Then, a fit is applied where the Gaussian is constrained to flatten at the final interference fringe. This provides a realistic value for 1/e radius, but does not track its profile exactly. A visual comparison of these fitting routines is given as Figure 6.12a).

	A		**B**	
Mirror Diameter	**X** (μm)	**Y** (μm)	**X** (μm)	**Y** (μm)
Phase Interferometry	58.03	67.42	88.66	97.86
Clipping Losses	51.65	62.34	98.00	98.00
Radius of Curvature	**X** (μm)	**Y** (μm)	**X** (μm)	**Y** (μm)
PSI - at σ_m	1064.86	861.80	1457	1765
TEM Splitting	1065.07	858.90	1419	1698
PSI - at centre	837.24	742.53	1080.40	1303.39
Clipping Losses	836.61	738.85	1079.95	-

Table 6.2: **Pyramidal Mirror Geometries.** A comparison of the geometrical properties measured for two pyramidal mirrors is given. These include measurements made by phase-scanning interferometry, the onset of clipping losses in length scanning and transverse mode splitting. Mirror curvature is grouped into two sections, demonstrating how a careful selection of the spherical fitting radius in interferometric reconstruction can give consistency with spectrally measured data. **A)** Properties of Pyramidal Mirror A. **B)** Properties of Pyramidal Mirror B. Here, mirror curvature could not be measured by clipping loss in one axis due to the complex instability profile seen in Figure 6.11.

Mirror curvature, as measured by transverse mode splitting, was in good agreement

with surface data when a spherical fit of the unconstrained Gaussian depression was made up to the 1/e radius of the constrained Gaussian depression. This is given as location B in Figure 6.12b). However, curvature measured by clipping loss onset was consistent to a fit made to the very centre of the depression, indicated by location A in the same figure. A summary of all geometrical properties for Pyramidal Mirror A are given in Table 6.2A. In using these conventions, consistency between interferometric and spectral analysis fell within 5μm, the approximate error of spectral measurements.

This accuracy was reduced for Pyramidal Mirror B, whose slightly non-Gaussian profile made it difficult to define a 1/e waist. Measurements of spectral curvature and interferometry were now within 50μm of each other, as shown in Table 6.2B. The minor-axis curvature measured by interferometry, $R_{oc} = 1080.40\mu$m, was in agreement with the smallest curvature seen in clipping, $R_{oc} = 1079.95\mu$m. However, curvature in the major axis, $R_{oc} \approx 1303\mu$m, is located within the complex instability region of clipping losses and therefore could not be verified. A more accurate determination may have been possible if coherence scanning interferometry had been used instead for surface reconstruction, but would have required greater time for manufacture.

6.3.4 Ellipticity and Birefringence

We regard mirrors to be of notable ellipticity if their fundamental cavity resonance is split into polarisation modes, where splitting is at least on order with linewidth. Accordingly, when the cavity addresses atomic structure, considerable polarisation dependent coupling will arise [38]. Repeating the result of [39], the polarisation splitting of the fundamental mode is given by:

$$\Delta\omega_{\varepsilon} = \frac{c\lambda_{probe}}{4\pi L_{cav}} \frac{R_{oc,H} - R_{oc,V}}{R_{oc,H} R_{oc,V}} = 2\pi\Delta\nu_{\varepsilon}, \tag{6.7}$$

where we have assumed that the mirror is orientated such that $R_{oc,H} > R_{oc,V}$. We compare this splitting to the minimum linewidth that a cavity can achieve in a prac-

tical implementation, given its inherent losses and maximum length. As can be seen from the form of the above equation, highly curved mirrors are more susceptible to modal splitting. Accordingly, of the geometries presented in Table 6.2, only A was considered as potentially birefringent, despite their ε values being similar, $\varepsilon_A = 0.44$ and $\varepsilon_B = 0.41$, given by their higher order mode splitting measurements.

An effective method of measuring fundamental mode splitting is to address the cavity with linearly polarised light, whose polarisation vector is orientated halfway between its major and minor axes of ellipticity. This will address both polarisation eigenmodes of the cavity, whose transmitted resonance profiles can be separated on a polarising beam-splitter and independently measured. When these two resonance profiles are overlapped in time, their respective detuning can be determined using standard sideband frequency markers. However, given that we did not have two APD detectors available, we were required to rely upon the joint detection of the eigenmodes.

To characterise the fundamental mode splitting of a pyramidal mirror, a half wave-plate was placed in front of the cavity, rotated in steps and the transmission signal monitored. The beamline for this measurement is given as Figure 6.2, configuration A. An example of the evolving resonance profile, for Pyramidal Mirror A, is given as Figure 6.13c). Two distinct polarisation modes can be seen at each wave-plate position, with a changing distribution of input power coupling.

To compare the properties of each polarisation eigenmode, the wave-plate was set to distribute power equally and a two-Lorentzian fit made. This deconstructed line-shape is given as Figure 6.13a). The linewidths of the constituent modes differed, corresponding to a finesse of 158,817 $\pm$ 2,425 in the vertical axis and 143,744 $\pm$ 3,781 in the horizontal. This suggests that in addition to geometrically induced birefringence, there exists birefringence in the dielectric stack also [122]. In this case, the dielectric layer spacing along the vertical axis is closer to $\lambda_{\text{probe}}/4$, giving a higher reflectivity. Further,

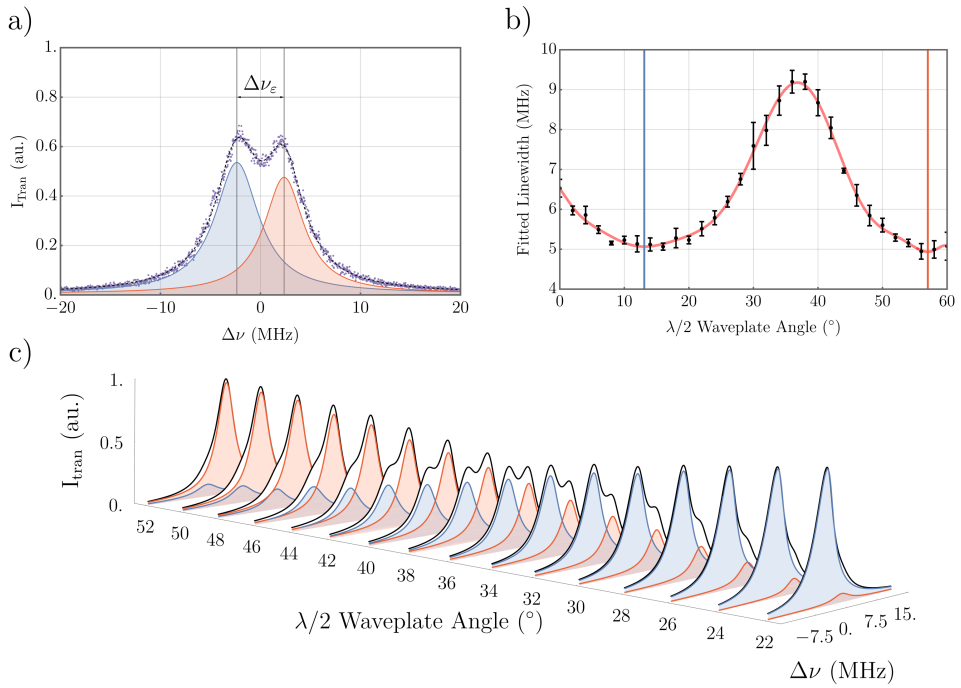

Figure 6.13: **Mirror Ellipticity. a)** A birefringent cavity, formed using Pyramidal Mirror A, is addressed with linearly polarised light aligned between its axes of ellipticity. This couples equal intensity into its two modes, with a frequency splitting given by $\Delta\nu_\varepsilon$. Captured data points are shown, with a joint Lorentzian fit in black an its decomposition in colour. **b)** As the probe beam polarisation approaches an eigenmode of the cavity, its orthogonal mode cannot be well resolved from the line-shape given in a), due to its low coupling. Instead, we fit a single Lorentzian function to the transmitted spectrum, where the narrowest linewidth corresponds to lowest excitation of the orthogonal mode. **c)** As an illustration of controlling coupling into each birefringent mode, the measurement made in a) is repeated over a $\pm 15°$ half waveplate rotation range. However, individual data points have not been plotted, only their fits.

the measured splitting was $\Delta\nu_\varepsilon = 4.72$MHz, which was greater than that the value of $\Delta\omega_\varepsilon = 3.4 \times 2\pi$MHz anticipated by mirror geometry. Where not otherwise stated, the cavity mirrors have been addressed with horizontally polarised light, potentially decreasing measured reflectivity.

When power was predominately coupled into a single polarisation mode, it became difficult to observe the orthogonal mode in their joint line-shape. To gain further insight of these regions, a single Lorentzian was fitted to the measured transmission profile over a 60° wave-plate rotation. The results of this characterisation are given as Figure 6.13b). The position of maximal linewidth corresponds well to the position of equal power distribution. Interestingly, however, the two linewidth minima are higher than those measured in 6.13a), which implies that linearly polarised light does not perfectly address the eigenmodes of the cavity and they have some elliptical phase.

Given the lower curvature of Pyramidal Mirror B, no birefringence could be observed at its maximum 'clipping free' length, when paired with HR reference mirrors. The minimum linewidth measured in this characterisation was $\Delta\omega_c = 1.80 \pm 0.13$MHz, which is similar to its expected splitting of $\Delta\omega_\varepsilon = 1.75 \times 2\pi$MHz. Given that the laser linewidth was 450kHz, this would be challenging to resolve. We are unlikely to pair this pyramidal mirror with a high reflectivity reference substrate[13] in any experimental setting, due to our requirement for finite cavity transmission. Thus, for cavity fabrication considered in the following chapter, this mirror is considered non-elliptical.

6.3.5 Review of Fabrication Process

Following the spectral analysis outlined in this section, our process of interferometric surface reconstruction is now able to anticipate the optical properties of pyramidal mirrors. In future manufacturing cycles, this will allow the prediction of cavity properties

[13] These mirrors have demonstrated a finesse of 1,900,000, considered one of the highest measurements observed at this wavelength [122].

at the time of mirror fabrication. Returning to the discussion of Chapter 4, our ablation process has demonstrated a desirable regime of mirror geometry, advantageous for strongly coupled interactions with trapped neutral atoms. However, we are also now in the position to consider potential improvements to our ablation method, chiefly by evaluation of the micro-patterning technique.

The micro-pattern of Pyramidal Mirror B favoured an expanded diameter, leaving the inherent ellipticity of the ablation laser uncompensated for. This produced a standard feature eccentricity of $\varepsilon \approx 0.4$, which was a beneficial trade-off in achieving a representative diameter of $\sigma_m \approx 100\mu$m. Due to modest surface scattering losses, millimetre scale R_{oc} and a requirement for transmission, implementation of a science cavity with $\Delta\omega_\varepsilon \geq \Delta\omega_c$ is unlikely with this mirror. However, with Pyramidal Mirror A a practical scenario was observed where $\Delta\omega_\varepsilon \approx \Delta\omega_c$. In this case eccentricity, $\varepsilon = 0.44$, was increased by the feature approaching the edge of the rectangular pyramidal facet and experiencing non-isotropic heat transport during ablation. Feature diameter was also limited by the small size of this facet, limiting the advantages of micro-patterning.

For the features fabricated on Pyramidal Mirror Array C, as shown in Figure 5.13, tailored ellipticity in the micro-pattern partially mitigated that of the ablation laser. This created features with $\varepsilon = 0.2$ and a standard diameter of $\sigma_m = 65\mu$m. A cavity formed using one of these features is accordingly less susceptible to birefringent behaviour than Pyramidal Mirror A. Its linewidth is also expected to be broader, as the onset of clipping loss occurs at a shorter mirror separation than Pyramidal Mirror B. What is apparent from this analysis is that the interplay of mirror diameter and ellipticity in our ablation technique must be considered in the context of cavity design. In this fabrication cycle we had notable flexibility with respect to birefringence, owing to a range of limitations in minimum cavity linewidth.

An improvement would be to reduce the ellipticity of the ablation laser, via a standard

implementation of spatial filtering, cylindrical lenses or anamorphic prism pair. This would allow the micro-pattern to have greater flexibility when low ellipticity and large diameter features are required. For the complete cancellation of ellipticity a substrate with cylindrical symmetry around the ablation site is desirable and the process of rotational smoothing [40] may be favoured. However, it may remain difficult to fully eliminate birefringence from the dielectric stack itself. The origin of this mechanism remains largely unexplained and is accordingly difficult to mitigate [126].

Following our discussion of Section 4.6.5, cavity fabrication may intend to enhance ellipticity, as to establish two fully resolved polarisation modes of the cavity. When creating a cavity from two identical elliptical mirrors, birefringence may be tailored via their respective rotation [39]. However, in this regard our pyramidal mirrors are limited, given their flat base for simplified mounting. As such, a continuation to micropatterning would be appropriate. We were limited in the fabrication of pyramidal mirrors by the number of base substrates available, but in principle a wide range of mirror ellipticity values may be achieved via this method [47].

6.4 Ring-down Spectroscopy

The data presented in this chapter indicates that the major source of incoherent loss in the pyramidal mirrors derives from 29ppm surface scattering and 8ppm absorption. The reduction of these losses will be a focus of further fabrication cycles, performed in consultation with the coating provider. For greater simplicity in this process, we wish to avoid characterising the properties of the dielectric coating in conjunction with the properties of the ablated feature. We have therefore developed a second apparatus for the measurement of mirror reflectivity, designed for operation with reference substrates which are coated alongside the pyramidal mirrors.

These reference mirrors would created using super-polished base substrates, with a ra-

dius of curvature commensurate with our standard specification of $R_{oc} = 50$mm. Since the scattering losses of these mirrors will be low, measurements of cavity finesse are expected to a direct measure of losses induced by the dielectric coating. Given the millimetre scale geometries of super-polished mirrors, we have a much greater freedom to tune cavity length during their characterisation. When compared to ablated mirrors, we are less limited by the onset of clipping losses and the resonator stability criterion.

We anticipate that these reference mirrors will have an increased reflectivity and therefore require measurement of a narrower linewidth. This will exacerbate some of the issues experienced in characterising the current pyramidal batch, including vibrational instabilities and convolution of the laser line-shape into the measured spectrum. This could be offset by decreasing the standard cavity length used, enhancing its linewidth, with a minimum limit given by the depth of the spherical feature. To overcome this limit, independent studies have chamfered the mirrors [35]. However, we wish to avoid post-processing of the mirrors after coating, to aid with simplicity and prevent induced damage. Therefore, we intend to characterise the reference substrates by the construction of a cavity ring-down spectrometer [79].

At extended cavity lengths, or when using higher reflectivity mirrors, the photonic lifetime of the cavity will surpass the time taken to cross resonance in standard length scanning spectroscopy. Therefore, the steady-state behaviour of the cavity will not be observed in its transmission, but one modified by a delay. For a cavity leaking light, the exponential decay of its transmission is characterised by a time constant, τ_{RD}, given by the following relationship:

$$\begin{aligned} I_T(t) &= I_0 e^{-\left(\frac{t}{\tau_{RD}}\right)}, \\ \text{where } \tau_{\mathrm{RD}} &= \frac{1}{c}\frac{L_{cav}}{1-\sqrt{R_1 R_2}}. \end{aligned} \tag{6.8}$$

For a cavity of known length, this time constant can be converted into mirror reflectivity,

thereby replacing the measurements of linewidth performed earlier. To measure this, we first bring the cavity onto resonance with the probe laser by slowly scanning its length and letting cavity transmission build to a given set-point. Then, laser illumination is rapidly switched off and the following exponential decay in output intensity measured. This method is fundamentally insensitive to environmental noise, as the measured decay profile does not depend on the changing resonance condition of the cavity. Indeed, the cavity must only be on resonance with the probe laser long enough to pass sufficient light into the cavity as to be detected in transmission. In contrast to our method of linewidth determination, a broad linewidth laser actually assists in this technique, as it can be more readily coupled into the cavity.

6.4.1 Overview of Apparatus

The system developed for cavity ring-down used a number of design features inherited from the apparatus of Section 6.1.2. However, it is a considerably more compact setup, with an effort made to simplify mirror exchange and cavity alignment. The novel elements necessary for ring-down have also been introduced, including the development of a fast optical switch and triggering method. A schematic for the device is given as Figure 6.14.

The spectrometer is primarily constructed from cage optomechanics, where the course alignment of each optical element is ensured by their mounting to a common set of rods. One of these plates, housing a planar mirror, is oscillated in the optical axis of the system using three piezoelectric chips. The opposing spherical mirror is contained within a five-axis kinematic mount, allowing for cavity alignment. This mount additionally contains a co-translating plano-singlet lens, for cavity mode-matching. This design was more robust to noise than the first apparatus discussed, as although it was subject to the same environmental fluctuations, the closed nature of the cage ensured that each mirror was vibrated in-phase. A large range of cavity length tuning was possible, by simply sliding the kinematic mount along the rods and occasionally updating the

focal length of the integrated lens. A CCD camera[14] is used for initial cavity alignment.

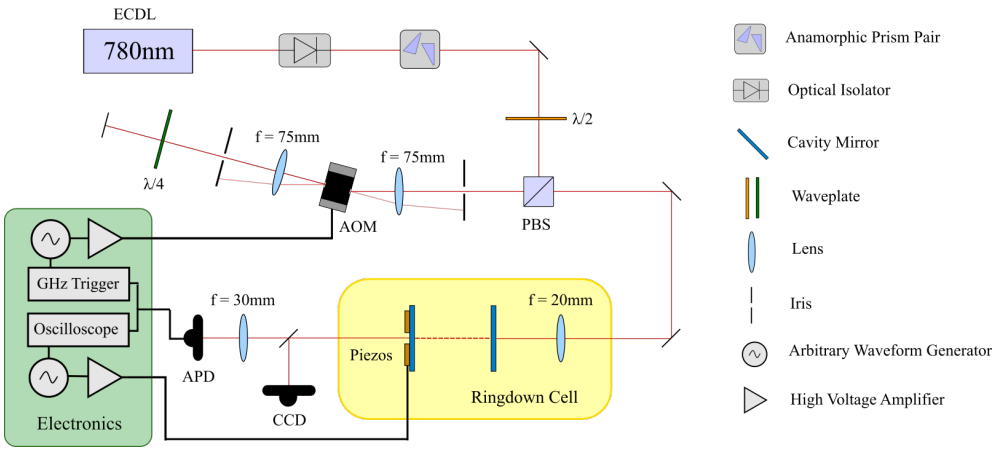

Figure 6.14: **Ring-down Spectrometer.** The beamline for the ring-down spectrometer is shown. A fast optical switch is created using an AOM in double-pass configuration. This switch diverts illumination of the cavity when a Schmitt trigger signal fires. This trigger occurs when cavity transmission, as measured on an avalanche photodiode, reaches a voltage set-point. The optical ring-down signal is then recorded on a high bandwidth oscilloscope.

The fast optical switch was realised by an acousto-optic modulator[15] (AOM) in double-pass configuration. Within an AOM is a piezoelectric transducer, used to generate sound waves within a sample of glass. These waves diffract incoming light, separating it into a spread of frequency shifted orders. We make use of the beam deflection applied by the AOM to rapidly switch light away from cavity illumination. In double-pass configuration, we retro-reflect the first order of the AOM output back into it, diffracting it for a second time and improving beam cancellation.

The piezoelectric element of the AOM required a 110MHz radio-frequency drive signal, supplied by an arbitrary waveform generator.[16] To observe cavity ring-down, this drive was rapidly switched off when the intra-cavity light intensity reached a reasonable level.

[14] Daheng Imaging MER-130-30UM-L
[15] Crystal Technologies AOMO 3110-120
[16] Keysight 33600A

This condition was monitored in cavity transmission by an APD.[17] When the voltage output of the APD exceeded a one volt set-point, the RF drive from the AWG would be cut, cancelling the deflection of the AOM light and de-coupling from the cavity within 33ns. To prevent triggering the optical switch on noise, a Schmitt trigger was developed, which would only fire when two voltage set-points were crossed. The voltage difference in these set-points was considerably larger than the standard noise amplitude, giving a reliable trigger action.

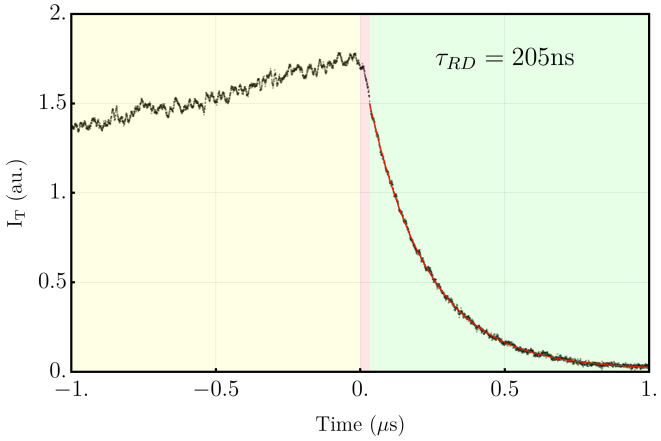

Figure 6.15: **Ring-down Characterisation.** An example ring-down measurement is given, broken into three shaded sections. In the first section, shown in yellow, cavity length is approaching peak resonance and light is building within it. In the second section, shown in red, the trigger set-point has been reached and there is a 33ns period in which the AOM switches light away from the cavity. In the final section, shown in green, the exponential decay of light in the cavity is fitted for its decay constant and shown as a red line.

6.4.2 Reference Substrate Testing

In our system, the transition time of the fast optical switch defines the maximum measurable decay rate. This establishes the limit on the minimum mirror reflectivity that can be determined, when the full cavity length range is used. This limit is far more prominent than for maximum mirror reflectivity, as extended ring-down profiles are relatively simple to measure. So to examine the performance of the constructed system, a range of lower-quality dielectric mirrors were evaluated, coated alongside the pyramidal

[17] Thorlabs APD110A/M

mirrors.

The majority of our tested mirrors were planar, used as the base substrate[18] in the fabrication of extended cavity arrays. However, we are naturally unable to characterise a pair of planar mirrors, owing to the accuracy of alignment it would require. Instead, they were each paired with an $R_{oc} = 25$mm mirror,[19] with the same applied coating. The ring-down of paired spherical mirrors indicated losses of 132 ± 15ppm each, resulting from surface scattering. Testing of the planar mirrors then indicated losses of 42 ± 5ppm each, where uncertainty is derived from the variable quality of the 6 analysed mirrors. An example ring-down measurement of a planar mirror is given as Figure 6.15.

The standard cavity length used was $L_{cav} \approx 5$mm and the standard finesse measured was $\mathcal{F} \approx 35{,}000$; corresponding to a ring-down time of $\tau_{RD} = 200$ns. However, the system remained able to resolve decay constants of as low as $\tau_{RD} = 100$ns. The quality of these mirrors places them well within the regime of analysis by cavity length scanning. However, if reflectivity is increased in future manufacture, ring-down will not experience a limit on minimum cavity linewidth. Measuring reflectivity was also far more simple via this method, as it did not require extended mirror alignment or suppression of vibrations. If $\sigma_{rms} < 1$Å super-polished substrates are used, a more direct measurement of coating absorption is possible. Therefore, it acts as a complementary technique to length scanning spectroscopy in the future development of low loss mirrors.

6.5 Multi-Feature Resonators

A key attribute of pyramidal mirrors is their ability to act as host to an array of ablation features. The construction of a multi-feature resonator with optical access embodies a range of desired experimental traits, building towards greater experimental simplicity and scalability. In this section, we further motivate the development of a multi-feature

[18] Laseroptik S-00018
[19] Laseroptik S-00128

high-finesse optical resonator, before reporting on its first demonstration.

6.5.1 Experimental Applications

The operation of just two cavities within an ultra-high vacuum environment would normally require duplication of the optics and electronics required for their length stabilisation. However, if these two cavities are formed by features on the same base substrate, then positional feedback applied to one will stabilise the other. The stabilisation of a cavity without direct laser illumination, or 'dark' stabilisation, provides an excellent environment for atomic interaction. The laser light normally used for locking will no longer modify atomic states via the AC Stark effect or change their occupation. This is normally only achieved by regular interruption of atom-resonant locking light, or shifting it off resonance by use of complex frequency references.[20]

A multi-cavity platform would also represent a notable advance in the implementation of scalable quantum information processing. If a re-configurable single atom dipole trap can be integrated into such a cavity, then multiple atom-cavity interfaces could be operated simultaneously. A critical aspect in the storage and transport of atoms by dipole trap is the length scale over which it is required [66]. In this sense, cavity modes formed on the same substrate are likely to be far more convenient.

6.5.2 Demonstration of Multi-Feature Co-Resonance

To evaluate the suitability of our pyramidal substrates as the host of multi-feature resonators, the four feature Pyramidal Mirror C was installed into the characterisation apparatus. This substrate comprised three $R_{\mathrm{oc}} = 650\mu\mathrm{m}$, $\sigma_m = 65\mu\mathrm{m}$, $\varepsilon = 0.2$ features and one $R_{\mathrm{oc}} = 250\mu\mathrm{m}$, $\sigma_m = 63\mu\mathrm{m}$, $\varepsilon = 0.13$ feature. The pyramidal facet had total dimensions of $380\mu\mathrm{m} \times 320\mu\mathrm{m}$, with $100\mu\mathrm{m}$ spacing between each feature. The probe laser beam-line was expanded to have four independently steerable beams, split using polarisation optics on the laser output and recombined before the cavity.

[20] We specifically refer to the use of transfer cavities and optical frequency combs.

All four probe beams were initially overlapped and coupling optimised into a single feature. Linear displacement was then made to three of the beams, coupling them into the remaining cavity modes. The position of the cavity mode on each feature was used to bring the planar reference mirror into a mutual alignment. Initially, the cavities resonated discordantly, with a clear time delay between their fundamental mode resonances when length was oscillated.

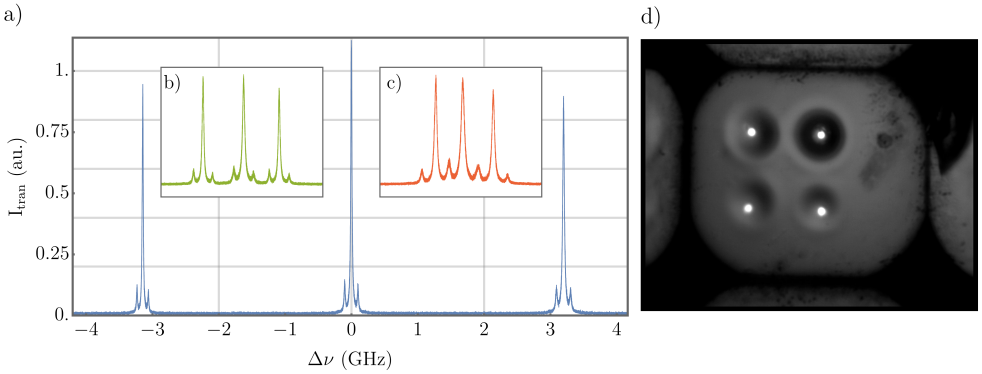

Figure 6.16: **Multi-Feature Resonator.** The joint transmission profile of three cavity features is shown, whose frequency separation is controlled by static piezoelectric rotations of a planar reference mirror. These resonances can be brought arbitrarily close together, where $\Delta\nu = 0$ denotes equal cavity length. To demonstrate this control, traces are shown for **a)** $\Delta\nu = 3.15$GHz **b)** $\Delta\nu = 415$MHz **c)** $\Delta\nu = 200$MHz. Separations lower than this would cause overlap of the features, making their illustration more difficult. **d)** An optical image of Pyramid C, corresponding to three-feature co-resonance on a fundamental mode. The high curvature feature in the top right is actually several GHz detuned here, but becomes resonant over the integration time of the camera.

To rectify the differences in cavity length, the reference mirror was tipped and tilted using piezoelectric actuators.[21] This allowed two features to be moved into co-resonance in the horizontal axis and two in the vertical axis. Overall, this allowed for three features to have their fundamental modes in simultaneous resonance during length scanning. An image was captured during co-resonance, given as Figure 6.16d). The detuning between

[21] Thorlabs MAX602D/M

each resonance could be arbitrarily tuned within a free-spectral range using the voltage applied to the piezoelectric elements, as demonstrated in Figure 6.16a-c). However, there was no degree of freedom in which to tune the length offset of the fourth mode, which remained approximately 5GHz detuned from the others. This was chosen to be the higher curvature feature, whose residual detuning can be explained by greater feature depth.

6.5.3 Towards Multiple Atom-Cavity Interfaces.

We have demonstrated a range of foundational elements in establishing a multi-feature cavity in ultra-high vacuum with single atom trapping. The spacing between features is commensurate with a large dipole trapping pattern generated by spatial light modulation and the facet size would allow for reasonable optical access. The ability to electronically tune the respective detuning between neighbouring cavity modes allows up to three co-resonant systems. The fourth feature could be used for their length stabilisation, where an inherent length offset can be compensated in the wavelength of the locking laser. If there are more than three features in an array, any differences in feature depth cannot be compensated by changing planar mirror angle. Therefore, to expand the possible number of co-resonant cavities, a very high precision would have to be reliably achieved in their ablation.

Chapter 7

Implemented Science Cavity

This chapter describes the design, construction and operation of an experimental system for strong atom-cavity coupling, using a cavity formed from tapered mirrors. Detailed discussions are given on the cavity alignment process, its vibrational stability and the ongoing integration of atomic cooling and trapping methods. Finally, we consider the future direction of this experimental platform and potential routes towards the creation of an elementary quantum network.

7.1 Overview

The greater goal of this work is to establish a technical platform for the development of cavity mediated quantum information processing. In the first instance, we intend to produce two implementations of a single trapped atom strongly interacting with a resonant cavity. These systems would act as a generator of entanglement, engineering a joint atomic state by interaction of their produced photons [23, 24].

To do so, several lines of previously distinct research and development in our group need to be combined. As a continuation of the work described thus far, we perform the fabrication and evaluation of an operational optical cavity built from tapered mirrors. In recognition of our goals for a highly replicable system, this is manufactured as a modular element using novel alignment methods. Furthermore, the cavity mounting system has been specifically developed to promote vibrational stability.

We then intend to integrate the cavity with our established work in atomic cooling

and re-configurable dipole trapping [66, 140]. This has been developed in a similar vacuum chamber as our optical cavity, to allow for their efficient unification. However, restrictions in optical access have required a modification to our standard implementation of a magneto-optical trap (MOT). Accordingly, the cavity mounting structure and implementation environment has been designed in partnership with a colleague,[1] who ensured compatibility with a dipole trap. This cavity design permits atomic loading using compact MOT geometries, whose ongoing development will be discussed.

7.2 Science Cavity Properties

The characterisation work performed in the preceding chapter has shown that the pyramidal mirrors suffer from incoherent losses of $\mathcal{L}_{A+S} \approx 37$ppm. When compared to their transmission of $T \approx 2$ppm, their use for photon extraction from the cavity is limited. In order to build an efficient atom-cavity platform, we therefore pair them with high-transmission super-polished mirrors in a semi-planar geometry. These are from the same fabrication batch as the out-coupling mirrors used in characterisation, with $R_{oc} = 50$mm, $T = 38.2$ppm and $\mathcal{L}_{A+S} = 2$ppm.

From an atomic coupling perspective, the drawback of a semi-planar geometry is that the cavity mode waist is located on the planar mirror. Accordingly, the position of maximum coupling strength is unreachable by the dipole trap. However, given the long Rayleigh length of the cavity mode, high cooperativity can be achieved anywhere along the optical axis of the cavity. This hybrid geometry has been externally recognised as having a range of further benefits in experimental simplicity [141].

To ensure that the linewidth of the cavity is suitable for targeted atomic interaction, it was tailored to be similar to a demonstrated system [56], as discussed in Section 4.5.2. This necessitated an extended cavity length, which additionally creates excellent op-

[1] Mark IJspeert

Property	Symbol	Planar	Pyramid	Unit
Cavity Length	L_{cav}	450		μm
Radius of Curvature	R_{oc}	50	1.5	mm
Mirror Diameter	σ_m	300	100	μm
Facet Diameter	σ_s	400	275	μm
Transmission	T	38.2	1.75	ppm
Mirror Losses	$\mathcal{L}_{A+S}$	2.2	37	ppm
Cavity Linewidth	$\Delta\omega_{\mathrm{c}}$	4.2		2π MHz
Free Spectral Range	$\Delta\omega_{\mathrm{fsr}}$	330		2π GHz
Optical Access	θ_{oa}	48		$^\circ$
Coupling Strength	g_0	11		2π MHz
Photonic Decay	κ	2.1		2π MHz
Atomic Decay	γ	3		2π MHz
Cooperativity	C	9.6		-
Finesse	$\mathcal{F}$	79,330		-

Table 7.1: **Science Cavity Design Parameters.** Using the properties of Pyramidal Mirror B measured in the previous chapter, and design principles guided by Chapter 4, shown are the design parameters of our constructed science cavity. Cavity length has been increased towards its maximum value, as defined by clipping loss, to enhance optical access and minimise cavity linewidth. Despite the semi-planar geometry, high cooperativity and atom-cavity coupling can be realised. Then, reasonable photon extraction is expected by using a super-polished out-coupling mirror, overcoming the enhanced scattering losses of the pyramidal mirrors.

tical access for dipole trapping. To prevent the onset of clipping loss, we have used Pyramidal Mirror B, whose spectral and geometric properties are summarised in Tables 6.1 and 6.2. Using these measurements, the predicted properties of our science cavity are given as Table 7.1. These properties are verified later in this chapter, as part of the cavity fabrication process.

Overall, our cavity design has realised high cooperativity, narrow linewidth, negligible birefringence, 50% photon extraction efficiency and ample optical access. While the metrics of performance are lower than could be achieved using paired micro-mirrors, they still constitute a highly versatile and valuable prototype for the next stage of research. Once a batch of pyramidal mirrors has been fabricated with a high transmis-

sion coating, we may return to a geometrically symmetric cavity design. In combination with the manufacturing improvements recommended in previous chapters, this is expected to reach the state of the art in the metrics of cavity performance outlined in Chapter 4.

7.3 Cavity Implementation Environment

The design of the operational environment for our cavity has been modified in favour of ensuring versatility for atomic cooling and trapping methods. This required building the entire atom-cavity system within a thin rectangular glass cell, to provide excellent optical access from five sides of the chamber, assisting formation of the magneto-optical trap, dipole trap and optical interaction with the cavity. However, this glass cell also placed a significant restriction of the size of the cavity system and the nature of its mounting. We shall therefore describe the schematic of the vacuum chamber first, before outlining the cavity module tailored for operation within it and the mounting structure used to support it.

7.3.1 Vacuum Chamber

The operation of a narrow linewidth optical cavity generally requires a vacuum environment, to ensure suppression of acoustic noise during length stabilisation and a low probability of mirror contamination. However, this can be readily achieved in modest sub-atmospheric pressures, in contrast to the demanding requirements of single atom trapping. The lifetime of single trapped atoms has been seen to be strongly dependent on background pressure, requiring our systems to operate at approximately 10^{-10}mBar [66]. Accordingly, our entire experimental apparatus is located within an ultra-high vacuum (UHV) environment, whose overall schematic is given as Figure 7.1.

The core scientific section of the vacuum chamber is contained within the glass cell,[2]

[2] A custom fabrication by ColdQuanta Inc., Part CQOX0001D.

comprising flat borosilicate glass walls and connected to a steel flange via a thermally insulating interface. The walls are coated with dielectric anti-reflection (AR) coatings, for optimised transmission at 780nm and 1064nm. These wavelengths correspond to the ^{87}Rb D2 line and our dipole trapping laser light respectively. The internal cross section of this cell is 21mm×21mm, delimiting the maximum size of our cavity assembly.

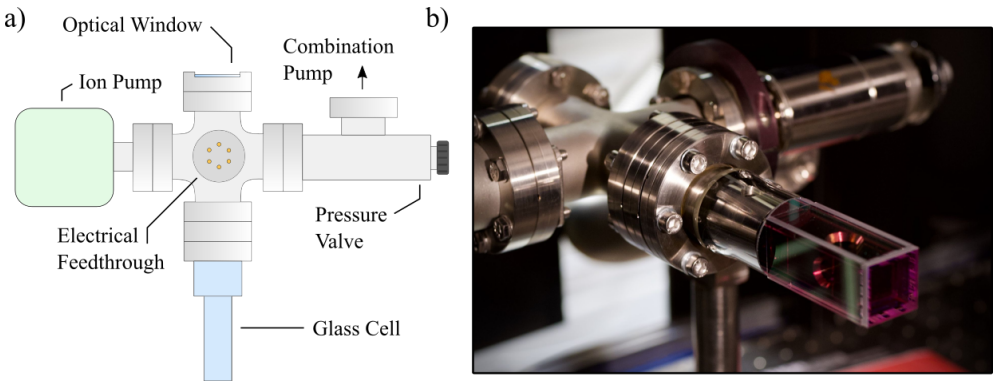

Figure 7.1: **Glass-Cell Vacuum Chamber. a)** A full schematic of the glass-cell vacuum chamber is given. An electrical feed-through is used for operation of the rubidium dispenser and length modulation of cavity piezos. An optical window allows illumination of the rear cavity mirror for length stabilisation. A combination turbo-molecular and roughing pump is used to reduce chamber pressure to 10^{-7}mBar, before being isolated by closure of the pressure valve and removed. The chamber is then maintained down to 10^{-10}mBar by operation of an ion pump. **b)** The cavity module, mounting structure and pyramid MOT have all been designed for integration within the compact glass cell. An image of this cell is shown, highlighted by its 780nm/1064nm AR coating for normal to 45° incidence. Inside, the cavity superstructure is visible, with its apertures for atomic cooling and trapping.

A steel mounting platform is installed into the cell, into which the cavity module and MOT optics are installed. The cell is attached to a five-way vacuum flange, which hosts an electrical feed-through and optical window. The feed-through is used to apply voltage modulation to the cavity piezoelectric transducers and to operate a resistively heated rubidium dispenser. The optical window is AR coated for 780nm and is used to illuminate the pyramidal cavity mirror. One arm of the five-way flange then leads to the ion pump[3] and integrated pressure sensor. The opposite arm leads to a combination

[3]Gamma Vacuum TiTan 25S-DI

turbo-molecular and roughing vacuum pump,[4] located behind a sealable UHV pressure valve.

In order to achieve the requisite vacuum pressure, the chamber was evacuated at 200°C, inside a purpose built oven. Chamber temperature was slowly increased to expel contaminants from its internal walls, removing them by the combination pumping stage. Once the pressure of the chamber became constant at this level, the ion pump was activated and the oven cooled. The valve was then sealed and the combination stage removed. At room temperature, the ion pump could reliably hold the chamber at 10^{-10}mBar without gas load.

Three vacuum chambers based on the above schematic were produced, to allow independent development of the core components of our experimental platform. The first chamber hosted an empty glass cell, for the development of atomic cooling and trapping. The next chamber contained only a prototype pyramid MOT, to allow for the study and optimisation of its cooling action. For the work of this thesis, a third chamber held all of the necessary elements for mounting and operation of the cavity, but did not require an ion pump for achieving pressures below 10^{-7}mBar. When each component reaches an adequate stage of maturity in their development, their unification into a single chamber should be a straightforward task.

7.3.2 Cavity Module

The optical cavity is based on our standard quasi-monolithic design [56], where the only degree of translational freedom comes in mounting each mirror onto a piezoelectric plate. These plates extended in the optical axis of the cavity, allowing for the electronic modulation of its length. Single-chip shear piezos were selected, as to allow a greater bandwidth of feedback than piezoelectric stacks.[5] Each piezo has a maximum

[4] Pfeiffer HiCube 80 Eco
[5] Noliac CSAP03

extension of 1.5μm, allowing traversal of several free spectral ranges. Piezo wires are coated with Kapton, due to its properties as malleable electrical insulator with excellent vacuum suitability.

In acknowledgement of the prototypical nature of this cavity, its mounting was designed with sufficient modularity as to allow for its simple removal and exchange. All cavity elements are secured by epoxy to a single metal plate, the design of which is illustrated in Figure 7.2. This plate was manufactured from non-magnetic 316LN stainless steel, to avoid perturbation of atomic structure by ferromagnetism.

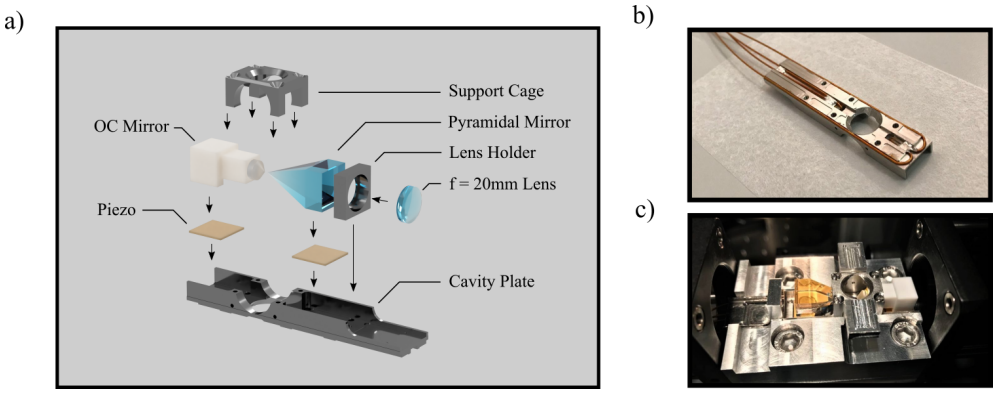

Figure 7.2: **Cavity Module. a)** An exploded rendering of the cavity module is given, showing its constituent parts. All elements are secured to the cavity plate, using vacuum compatible epoxy. At the front of the module is the coned semi-planar out-coupling (OC) mirror, held in a Macor mount. Due to its flat base, the pyramidal mirror is secured directly to the piezo. A support cage is placed over the mirrors, to improve the strength of the plate given its large dipole trapping apertures. A lens is integrated into the system, to aid mode-matching to the cavity via the pyramidal mirror. **b)** The underside of the plate contains indentations into which the piezo wires are secured, to ensure they do not obstruct the beamline in vacuum and aid its installation. **c)** The constructed cavity is shown in its alignment mount, before integration into vacuum.

The surface of this plate was highly structured to assist with cavity fabrication. This included indentations to ensure the correct alignment of the piezoelectric chips and dedicated epoxy wells to control its spread during curing. To aid with vacuum suitabil-

ity, several screw vents were installed and the plate was raised on runners. Dedicated channels were drilled on the underside of the plate to guide the piezo wires, to ensure they did not obscure the optical beamline of the cavity. Indentations were made for a beam splitter used in cavity alignment. Finally, a support bracket was attached over the dipole trapping aperture, to improve the strength of the structure.

This mounting system was originally designed for a paired pyramidal cavity, requiring modifications to be made to the planar out-coupling mirrors. These mirrors were first coned down to a 400μm facet diameter, as to be similar to the pyramidal mirrors. Then, to emulate the larger pyramidal geometry, they were secured within tapered mounts. These mounts were manufactured from Macor,[6] due to its machinability, high thermal tolerance and low out-gassing.

Given the length of the vacuum chamber and small size of cavity mode waist, it would be difficult to achieve perfect mode-matching into the pyramidal mirror using only external beam shaping elements. To reduce the divergence of an input beam, an $f = 20$mm plano-convex lens was integrated into the cavity module, focused onto the mode waist. Illumination of the cavity via the pyramidal mirror then facilitates length stabilisation. Generated photon capture could then be performed efficiently via the out-coupling mirror, which is sufficiently close to the edge of the glass cell to use external optics.

7.3.3 Mounting Superstructure

The design of the mounting structure used to support the cavity in vacuum was guided by post-analysis of two previous DPhil research projects in our group [116, 142]. These used fibre-tip mirrors, produced in collaboration with Jakob Reichel,[7] to build towards a strongly coupled atom-cavity system. However, this work was ultimately suspended, owing to a range of technical challenges. The issue of greatest applicability here is the

[6] Macor is a borosilicate glass-ceramic designed and manufactured by Corning Inc.
[7] Laboratoire Kastler Brossel, Paris.

low vibrational stability of the constructed cavity, which prevented its stabilisation on atomic resonance.

While the observation of vibrational noise in the measured spectrum of an optical cavity is common at atmospheric pressure, noise observed with the fibre-tip cavity persisted under vacuum. Electronic stabilisation of the cavity proved ineffective, making sustained atomic resonance impossible. The noise spectrum of the cavity was characterised and found to be highly similar to mechanical resonances predicted by finite element analysis of its mounting structure.[8] When piezoelectric feedback was applied over this set of frequencies, the fibres would simply oscillate in an uncontrollable manner. These mechanical resonances were too broad and numerous to simply apply an electronic filter to the error signal and the cavity too vibrationally sensitive to only apply low-bandwidth corrections.

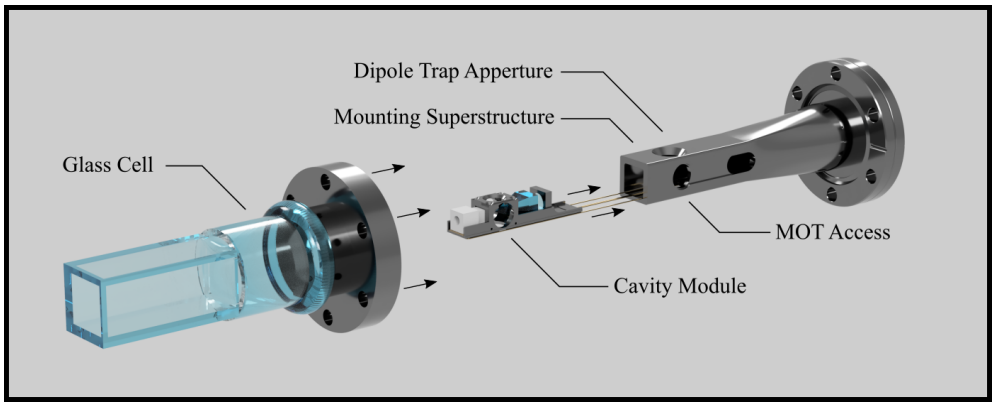

Figure 7.3: **Cavity Mounting Superstructure.** An exploded render is shown for the glass cell, cavity module and its mounting superstructure. The superstructure was designed to have the maximum mass possible to suppress vibration, limited by the internal dimensions of the cell and size of the pyramidal mirror bases. Multiple openings were installed into the superstructure, to allow bi-directional optical access to the cavity mode, to allow formation and imaging of the dipole trap and for bi-directional illumination of the pyramidal MOT. A further horizontal opening was included to clean rubidium from the integrated cavity lens by UV stimulated desorption.

[8]Fibre-tip cavity fabrication was performed by Marwan Mohammed, noise characterisation by Ezra Kassa and finite element analysis by Mark IJspeert.

This project therefore endeavoured to build a cavity system with significantly greater mechanical stability. This was one of the motivations for replacing the use of fibres in favour of tapered glass substrates, as they have a mass that prevents cantilever-like oscillation. Then, the cavity mounting structure was redesigned to eliminate strong mechanical resonances and accommodate the additional optics required for pyramidal mirror coupling. The revised cavity superstructure is shown in Figure 7.3. It consists of a hollow square tube, with outer dimensions that approach the boundary of the glass cell and internal dimensions which are slightly larger than the pyramidal mirror bases. This structure was fabricated from 316LN stainless steel.

The superstructure is designed such that an optical cavity module can be easily slid into position and secured with six base screws. The cavity may be illuminated from the rear of the cell, allowing interrogation of both cavity transmission and back reflection signals. An aperture is placed alongside the integrated cavity-coupling lens, to allow for its illumination by UV light, causing the removal of rubidium from its surface. Indeed, this UV stimulated atomic desorption is also useful for atomic trapping, as it provides a more rapid change of the rubidium pressure than the dispenser.

The mounting structure contains two vertical apertures for the injection and imaging of the dipole trapping beam. In the axis orthogonal to the dipole trap and cavity, a compact optical structure for establishing a MOT can be installed. This may be illuminated from both sides, to create a cold atomic flux for dipole trap loading and optical molasses at the position of the cavity waist.

7.4 Cavity Construction and Implementation

7.4.1 Alignment Process

To address the challenge of creating a well-aligned, highly compact and UHV compatible cavity module, a novel method of cavity fabrication was developed. The foremost

principle of our technique is to ensure that the laser beam used to probe cavity resonance during its construction represents the ideal position of the cavity mode when integrated into vacuum. Therefore, when we align each mirror to this beam, we create a cavity that is perfectly aligned with respect to our atomic cooling and trapping methods, as well as being readily laser-coupled for length stabilisation and output photon capture.

The technique proceeds in three steps. First, the probe laser is aligned to the cavity plate, setting the desired position of the cavity mode. Then, the planar mirror is installed into this beam and aligned for its perfect back reflection. The position of this mirror is secured using epoxy and the pyramidal mirror introduced to the system. This mirror is aligned to ensure that the cavity mode forms co-linear to the probe beam and its spectral properties are verified. The pyramidal mirror is then secured in position, completing alignment of the cavity.

A novel aspect of this method is that during their respective alignment processes, the position of each mirror is held fixed. Alignment is instead performed by co-translation of the probe beam, cavity plate, mode matching optics and detection apparatus as a monolithic structure. We will now discuss each step in detail, highlighting the advantages of this approach. A visual representation of this method is given as Figure 7.4.

Step 1: U-Bench Alignment

The optomechanical device developed for alignment is based on a U-Bench structure,[9] used to support all of the translating optical elements. Each side of the bench hosts a fibre coupler for injection and extraction of the 780nm probe laser beam. A bespoke alignment mount is installed into the centre of the bench, which positions the cavity plate in the correct position for attaching the mirrors. When this plate is installed into vacuum, the alignment mount ensures that the cavity mode will propagate through the

[9] Thorlabs CBB1/M

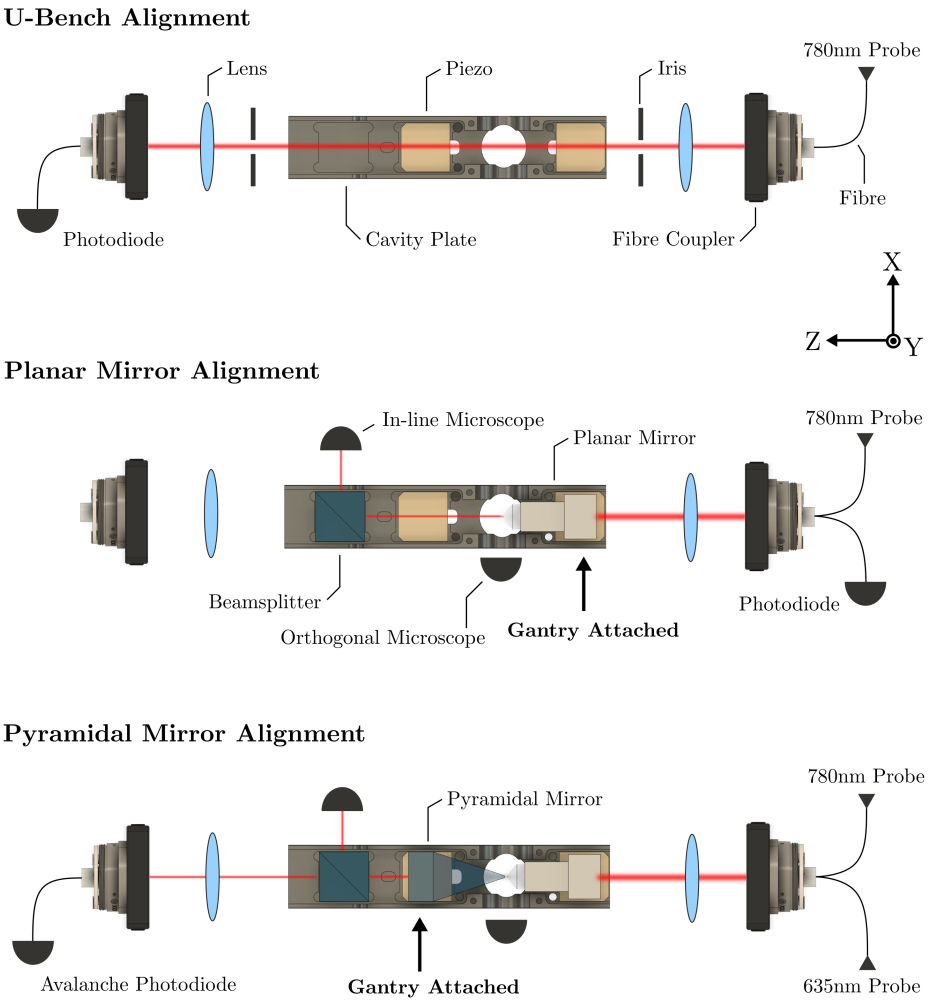

Figure 7.4: **Cavity Alignment Process.** The cavity alignment process comprises three key steps: alignment of the U-Bench, alignment of the planar mirror and alignment of the pyramidal mirror. A detailed procedure for each step is given in Section 7.4.1. The coordinate axes for this diagram are given in the lower right of the U-Bench alignment step.

centre of the glass cell.

The first step in this process is to align the probe beam to the central axis of the U-Bench. This is performed with kinematic elements in each fibre coupler, using a series of apertures placed along the bench. A photodiode[10] is attached to the output coupler, to provide an electronic measure of alignment.

Step 2: Planar Mirror Alignment

The super-polished planar mirror is installed first, by lowering it into the U-Bench on the arm of a gantry system. Its position in the propagation axis of the probe beam (Z axis in the associated figure), is set using apertures milled into the alignment mount. The position of these apertures matches the expected direction of cold atomic flux when installed into vacuum, for loading the intra-cavity dipole trap. The position of the mirror orthogonal to the probe beam (X and Y axes in the associated figure) is set using the installed beam splitter and *in-line* microscope. This images the facet of the mirror and is used to ensure that the probe beam is incident on its centre. Finally, the pitch and yaw of the U-Bench is tailored to ensure perfect back reflection of the probe beam from the mirror, re-coupling light into the fibre. The input of the fibre coupler hosts a fibreised beam splitter[11] and photodiode, allowing the extent of back reflection to be monitored electronically.

Each signature of alignment is optimised iteratively, until all are achieved simultaneously. Once in the correct position, the mirror is vertically retracted using the gantry arm. Epoxy is injected onto the piezo and the mirror returned to position. The system is illuminated with UV light[12] to cure the epoxy, at which point it may be unattached from the gantry. It then co-translates with the cavity plate and probe beam.

[10] Thorlabs PDA36A-EC
[11] Newport F-CPL-S22785-FCAPC
[12] Thorlabs CS2010

Step 3: Pyramidal Mirror Alignment

Next, the pyramidal mirror is introduced into the U-Bench on the gantry arm. Its position in the propagation axis of the probe beam (Z axis in the associated figure) sets cavity length, tailored using a calibrated *orthogonal* microscope. Cavity length is then rapidly scanned over a free spectral range, using the planar mirror piezo. This generates a cavity resonance signal on transmission, observed using the in-line microscope. The pitch and yaw of the U-Bench is altered to migrate the position of this mode to the centre of the pyramidal ablation feature. The transmission signal is also measured using an avalanche photodiode[13] (APD), placed on the fibreised output of the U-Bench. The position of the pyramidal mirror orthogonal to the probe beam (X and Y axes in the associated figure) is tailored to optimise the laser power coupled into the fundamental cavity mode, as measured using the APD.

As in the previous step, all signatures of alignment are optimised iteratively. If the formation of a cavity mode is lost entirely in this process, it may be restored by verifying the position of the 635nm probe beam. This is overlapped with the 780nm light on the input fibreised beam splitter, but unlike the 780nm light, is measurably transmitted by the mirrors when the cavity is off-resonance. Once the cavity mode is formed in the optimal position, the spectral properties of the cavity are determined by length scanning spectroscopy. The free spectral range of the cavity acts as verification of its length. Then, a measurement of linewidth is used to determine mirror losses and ensure no surface contamination has occurred. Finally, the pyramidal mirror is secured in position using epoxy, employing the same method as for the planar mirror.

In this final step some key advantages of the method are realised. Generally, measuring the spectral properties of a cavity during its alignment is challenging. The suppression of vibrational noise applied to cavity length normally requires secure mirror mount-

[13] Thorlabs APD110A/M

ing, thus restricting their translational freedom. In our method the mirrors are held fully stationary during their respective alignment steps, using a highly stable gantry structure. If this structure was required to have six axes of translation, ensuring stability would be considerably more difficult. Then, in co-translating the probe beam, mode-matching lenses and planar mirror, cavity alignment is performed rapidly with a single translation stage.[14] This is an improvement on the work of Chapter 6, in which the alignment of the cavity and steering of the probe beam were independent. It is now more similar to the alignment process of fibre-tip mirrors, where illumination and mode-matching optics inherently migrate with mirror position.

7.4.2 Alignment Apparatus

The aluminium U-Bench system is largely constructed from cage optomechanics, with a rendering given as Figure 7.5. This uses several smooth steel rods to guide the position of installed optical elements into a common propagation axis. Each side of the U-Bench hosts a five-axis kinematic fibre coupler,[15] for the injection and extraction of light. Each fibre-coupler has an associated lens in a kinematic xy mount, for cavity mode-matching. The bespoke aluminium alignment mount, hosting the cavity plate, is secured at the centre of the bench. For imaging the mirrors with an in-line 12x microscope,[16] a 10mm beam splitter is placed onto the cavity plate. The position of the beam splitter is guided by a square slot milled into the plate. The alignment mount has an aperture running orthogonal to the axis of the probe beam, to indicate the ideal position of the cavity mode for atomic loading in vacuum. An orthogonal 6x microscope images through this aperture, for setting mirror position. The entire U-Bench system is installed on a six-axis *NanoMax* translation stage, using a 3D printed adaptor.[17] This adaptor additionally acts as a support structure for the outermost optical elements, preventing their misalignment under gravity.

[14] Thorlabs MAX602D/M
[15] Thorlabs PAF-X-2-B
[16] Navitar 1-50487 with attachments.
[17] Using Formlabs Form 2 & Grey Pro Resin.

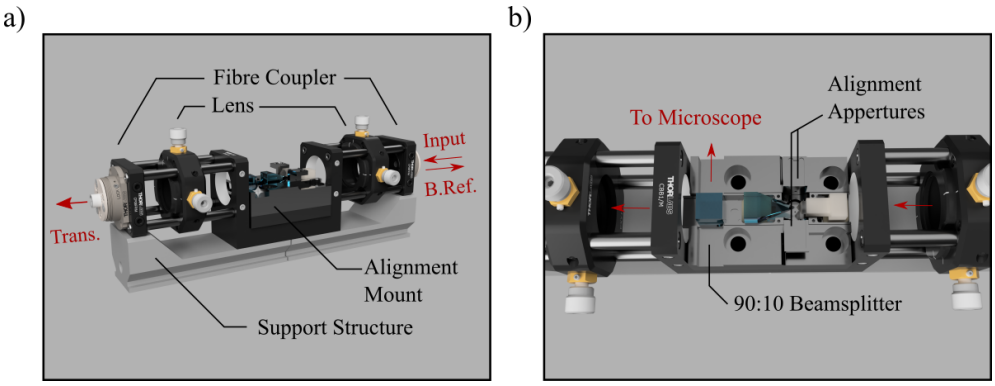

Figure 7.5: **U-Bench Cavity Alignment. a)** A side-view of the U-Bench is shown, contained within a 3D printed support structure for mounting to the translation stage. On each side of the bench are fibre couplers, allowing injection of the probe beam and capture of the back reflected and transmitted cavity signal. These are initially aligned to give perfect propagation of the probe beam through the centre of the U-Bench and are then automatically translated with it. Two lenses are installed into kinematic xy mounts, for fine tuning of cavity mode-matching. The alignment mount, holding the cavity plate to the ideal height, is installed into the U-Bench. **b)** This alignment mount contains an indentation for a beam splitter, to image the cavity mode by microscope. Alignment apertures are placed orthogonal to the cavity axis, to allow its mode to be aligned to the intended position of atomic flux and dipole trapping.

During their alignment, each tapered mirror is held over the cavity plate by the gantry system[18] illustrated in Figure 7.6. This gantry system is constructed from 66mm thick aluminium posts, with dovetail rails along all sides. Right-angle clamps attach to these rails, connecting the posts together. Each mirror is secured into an interchangeable PEEK[19] claw using a teflon coated setscrew. This claw is attached to a stainless steel arm, in turn attached to the horizontal gantry post. This arm may be moved in each Cartesian axis, by sliding the right angle clamps along the rails and re-securing them. Fine tuning of mirror position in the vertical axis is provided by a linear micrometer stage.[20] This is used for the controlled lowering and retraction of each mirror, during

[18]This design was based on the standard protocol for the movement of shipping containers.
[19]Polyether ether ketone (PEEK) is an engineering polymer used for its high machining tolerance and relatively soft contact with an installed mirror.
[20]Thorlabs MS1S/M

their initial installation and injection of epoxy.

7.4.3 Science Cavity Fabrication Process

The U-Bench alignment method was used in the production of our tapered science cavity, whose anticipated spectral properties were given as Table 7.1. Its full fabrication process shall now be described, including assembly of the cavity module and its installation into ultra-high vacuum.

For UHV suitability, it was crucial that surface contamination of the cavity module was minimised. To this effort, each non-optical component was ultrasonically washed in aqueous detergent, acetone, isopropanol and methanol before use. The cavity mirrors and integrated mode-matching lens were cleaned using First Contact instead, to minimise residual solvent. All fabrication was performed inside a lamina airflow enclosure with HEPA filtered input. This minimised the presence of dust that could have adhered to the cavity mirrors.

The first stage of cavity module assembly was to attach the shear piezos to the cavity plate. Their position was guided by slots milled into the plate and verified under microscope. They were attached using electrically conductive Epo-Tek H20E epoxy, injected into dedicated wells. This epoxy was selected to ensure a good electrical contact between the cavity plate and grounded side of the piezo. It also has considerably greater viscosity than standard UHV epoxies, preventing its creep during curing. If the spread of epoxy is not well controlled in cavity fabrication, it is known to flow over the edges of piezos and seize their extension. The piezo was tested for electrical shorting by measurement of its capacitance and for elementary operation via the generation of acoustic waves when driven at kilohertz frequencies.

The cavity wires were bent into their dedicated channels on the base of the plate and secured using small amounts of thermally curing Epo-Tek 353ND epoxy. The super-

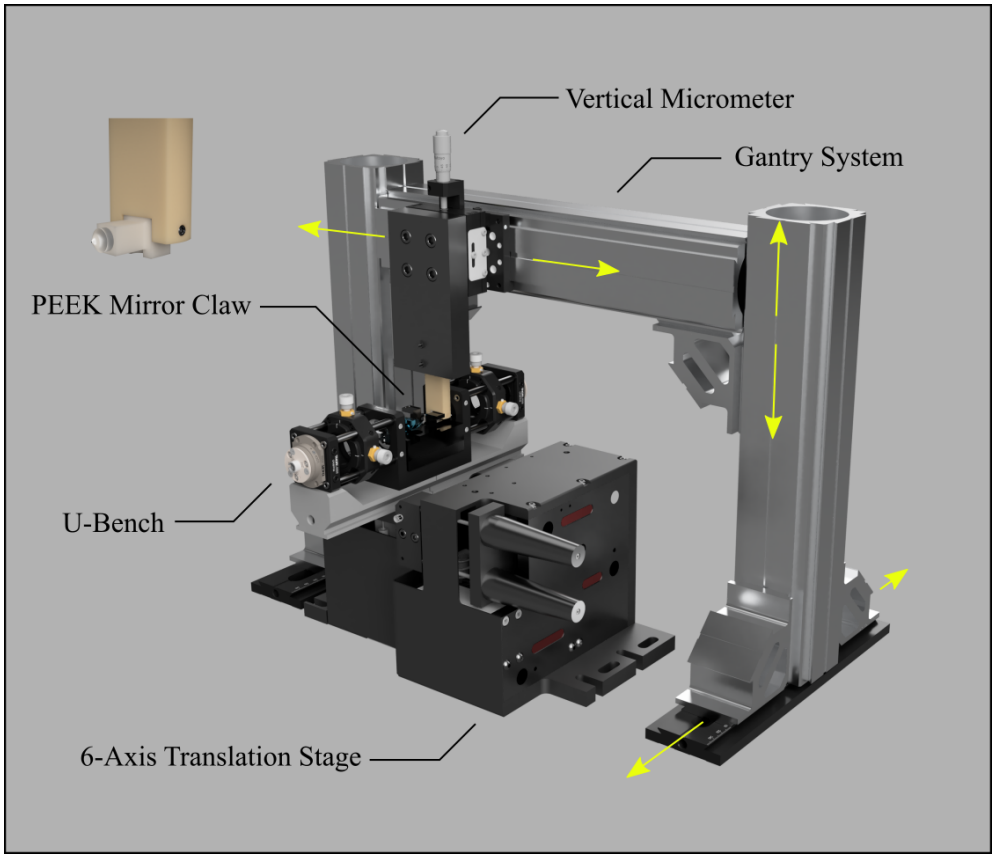

Figure 7.6: **Gantry Alignment System.** A full schematic of the aluminium gantry is shown, positioned for alignment of the first cavity mirror. The mirror is held by a PEEK claw, shown as the top-left inset. The height of this mirror can be fine-tuned by micrometer, allowing its controlled retraction for the injection of epoxy on the piezo. This claw can be coarsely translated in all Cartesian axes, shown by yellow arrows, by sliding and re-securing its mounts along the gantry system rails. This is used to bring the mirror claw into an ideal position for the installation of each cavity element. However, all cavity alignment is performed by translation of the *NanoMax* translation stage, with the attached U-Bench.

polished mirror and integrated cavity mode-matching lens were also secured into their respective mounts by 353ND epoxy. However, thermally curing epoxy is unsuitable for alignment critical components, such as securing the cavity mirrors in position. The heat used in curing would cause the mirrors to shift position, resulting in misalignment of the final cavity mode.

For securing the cavity mirrors in place, UV curing epoxies are favourable. Standard examples include Epo-Tek OG116-31 and OG198-54, with the latter having the option to thermally post-cure areas of poor UV exposure. However, owing to manufacturing delays, the previously untested Norland Optical Adhesive 88 was used instead. This reports low out-gassing for vacuum suitability, however to ensure its minimisation, the epoxy was first pumped down to 10^{-5}mBar in a constructed vacuum chamber. This chamber hosted several optical windows, which allowed observation of large quantities of volatile components in the epoxy boiling off at low pressures.

The planar mirror was aligned and secured to the cavity plate as described in the previous section. Before the pyramidal mirror was secured by epoxy, the support cage was installed onto the cavity plate. This prevented any induced flexure of the plate modifying cavity alignment. After curing, the integrated lens was installed using the gantry arm. Its position was tailored by removing the cavity output mode-matching lens in the U-Bench and re-optimising the cavity transmission signal as observed on the avalanche photodiode. The integrated lens was also secured by UV curing epoxy.

The completed cavity module was removed from the alignment mount and installed into the mounting superstructure within the vacuum chamber. The chamber was evacuated by turbo-molecular pump and the properties of the cavity re-determined by length scanning spectroscopy. These values are given as Table 7.2, which are highly consistent with those predicted in our design. However, the cavity is $\approx 25\mu$m longer than intended, corresponding to the resolution of the microscope used in alignment. The small increase

Property	Symbol	Design	Measured	Unit
Length	L_{cav}	450	476.18 ± 3.8	μm
Linewidth	$\Delta\omega_{\mathrm{c}}$	4.2	4.17 ± 0.13	2π MHz
Free Spectral Range	$\Delta\omega_{\mathrm{fsr}}$	330	315.0 ± 0.3	2π GHz
Pyramid Losses	$\mathcal{L}_T$	37.3	40.8 ± 1.2	ppm
Finesse	$\mathcal{F}$	79,330	75,904 $\pm$ 2657	-

Table 7.2: **Measured Science Cavity Properties.** After being installed into vacuum, the spectral properties of the tapered mirror cavity were redetermined. This indicated a small increase in cavity length and thus the initial onset of clipping loss. However, owing to the margin of error applied to cavity design, this only induced a 3ppm additional loss to the cavity. Overall, these properties compare favourably to the intended design properties given in Table 7.1. Errors are determined statistically.

in clipping losses at this length broadened the cavity linewidth. However, since this was perfectly, and somewhat coincidentally, offset by the longer cavity length, it was not noticed during spectral interrogation of the cavity during fabrication. However, the additional losses of 3ppm do not notably affect cavity performance.

7.4.4 Frequency Stabilisation

For coherent manipulation of an atomic state, the resonance frequency of the optical cavity must remain at a fixed detuning from its targeted transition. For a cavity finesse in the regime of 10^5, this implies a required length stability in the order of picometres. This is exceptionally challenging to achieve in a standard laboratory environment, given the acoustic, thermal and vibrational noise present. Instead, cavity length is stabilised by electronic feedback, applied via piezoelectric transducers. To create an optical measure of cavity length suitable for the generation of an error signal, its resonance with a frequency stable laser is probed.

Ideally, the laser frequency used for length stabilisation is sufficiency far detuned from any relevant atomic transition as to avoid interference with the process under study. This requires the laser to address a longitudinal cavity mode that is distinct from the

mode used for atomic interaction, separated in frequency by an integer number of free spectral ranges. However, laser frequency must also be within the stopband of the dielectric coating, as to ensure a narrow cavity linewidth. This usually requires a stable frequency reference with inherent tunability, with examples including an optical frequency comb, a transfer cavity or the use of electro-optic and acousto-optic modulators. If a suitable device is unavailable, the cavity may be stabilised with the exact laser frequency of its intended atomic resonance, with illumination interrupted during atomic interaction. This is a sample and hold technique, requiring excellent passive stability of the cavity over the interaction period [120].

In future experimentation, we intend to stabilise our cavity detuned from atomic resonance using an optical frequency comb.[21] This is currently installed within our laboratory for locking a paired super-polished mirror cavity [56]. The extension of this locking scheme to support two cavities simultaneously is ongoing. In the current stage of development, it is sufficient to characterise the cavity by stabilising it to an atom-resonant laser. Since this laser frequency also matches the design wavelength of the dielectric stack, it experiences a maximum cavity finesse. Therefore, the cavity resonance signal is maximally susceptible to vibrational noise. Locking at this wavelength ensures that the cavity has sufficient vibrational stability to be locked to any practical laser frequency in future.

Cavity locking light is produced by a Toptica DL Pro external cavity diode laser (ECDL), tuned to the $|F = 2\rangle \rightarrow C_{13}$ crossover line (780.246nm) in saturated Doppler-free absorption spectroscopy of ^{87}Rb. Locking the laser to this spectral feature is performed by the Pound-Drever-Hall (PDH) technique [143, 144]. The required 25MHz frequency sidebands are generated by RF modulation of laser current, with demodulation, signal processing and PID locking performed with a Toptica DLC Pro controller. The application of sidebands and their power distribution is monitored on a scanning

[21] Menlo Systems FC1500-250-WG

Fabry-Perót interferometer, with the central frequency of the laser verified by wavelength meter.[22]

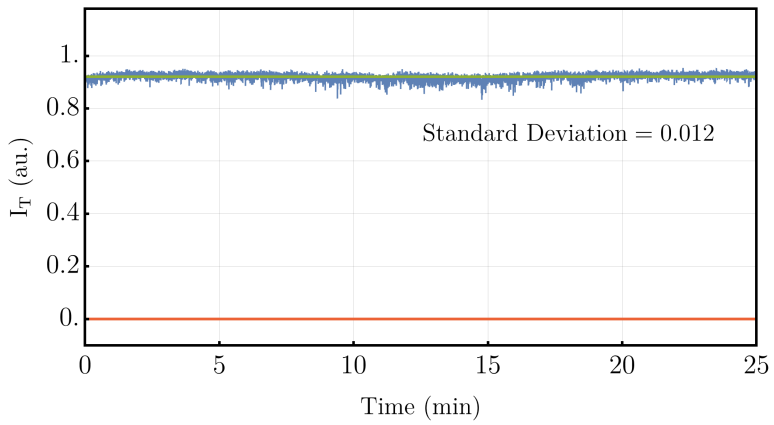

Figure 7.7: **Long-Term Cavity Stability.** The transmission of frequency stable laser light by the cavity is shown for a period of several minutes, with PDH stabilisation of length performed. Mean cavity transmission is indicated by a green line and off-resonant transmission by a red line. Fluctuations in cavity transmission can be interpreted as a changing resonance frequency of the cavity, with respect to the probe laser, which is not corrected by the feedback electronics. These deviations could be significantly reduced with the application of higher frequency feedback, in a range of greater than 1kHz. However, the current performance of the system is suitable for engineering persistent atom-cavity interactions.

The science cavity is locked on its transmission signal, measured using an avalanche photodiode.[23] This uses the PDH method, with demodulation, filtering and PID locking performed using a second channel of the DLC Pro controller. The bandwidth of feedback generated by this controller is limited to ≈1kHz, applied to the planar mirror piezo of the cavity. The pyramidal mirror piezo is used to apply a static offset to cavity length, by driving it with a fixed voltage.[24] This is used to centre the cavity length on resonance, against the action of slow thermal drifts. This increases the duration of time that the cavity remains locked, as it alleviates the output voltage limit of the feedback electronics.

[22] HighFinesse WS6
[23] Thorlabs APD110A/M
[24] Keysight 33600A & Falco Systems WMA-01LF

Under the current scheme the cavity remains locked for several hours, with disruption normally occurring only as a result of severe acoustic disturbance. An example of long-term cavity behaviour is given as Figure 7.7. There is excellent scope to improve this stability, chiefly by increasing the relatively low bandwidth of feedback. This feedback may also be split in frequency, applying high frequency corrections to one piezo and low frequency to the other. This will assist in tracking thermal drifts without manual intervention. However, current performance demonstrates that the vibrational stability of the cavity is sufficient for engineering persistent atomic interactions. This has been a notable challenge in the implementation of fibre-tip cavities.

In future work the vibrational stability of the cavity will be probed quantitatively, by examining the mechanical resonances of its structure, the response of the cavity to applied feedback and the isolation that its implementation environment provides. This will allow feedback to the design of the cavity superstructure and direct comparison with similar studies performed on fibre-tip cavities, which were specifically developed for passive stability [49].

7.5 Towards Intra-Cavity Atom Trapping

Atomic cooling and trapping has been a parallel line of research to optical cavity development within our research group. This is discussed in a selection of works [65, 66, 140], attributed to the authors therein. However, to integrate these methods into our intended cavity platform, significant revisions and extension will be required. This section summarises our goals for intra-cavity single atom trapping, past achievements towards this effort and the outstanding requirements.

7.5.1 Dipole Trapping

Many studies of cavity quantum electrodynamics in the optical domain have used single non-localised atoms, traversing the cavity volume with a slow velocity [19]. An effective method of creating this emitter stream is via an atomic fountain [145]. This requires the formation of a magneto-optical trap (MOT), generating a cloud of cold atoms. The laser beams used to create the MOT are then variably detuned, propelling atoms towards the cavity mode. If the density of the cloud and ballistic properties of the fountain are correctly tailored, long interaction times can be realised as a single atom traverses the mode. With a high atom-cavity coupling strength, this interaction period can be sufficient to allow multiple coherent operations on the atom, including several successive photon emission events [56].

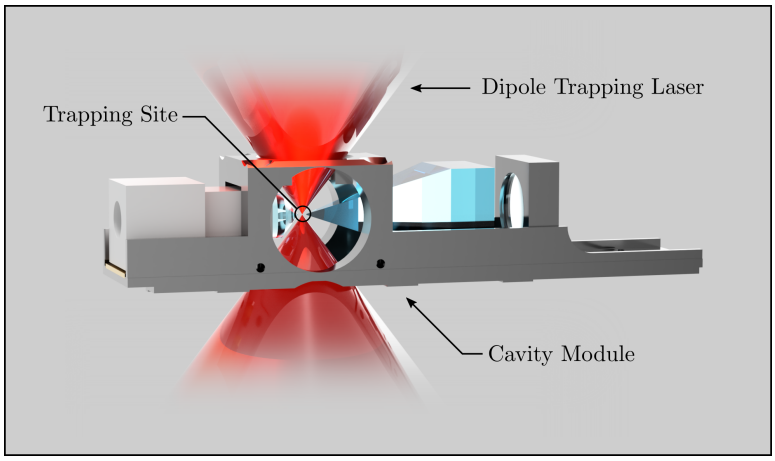

Figure 7.8: **Anticipated Intra-Cavity Dipole Trap.** A rendering of the cavity module is shown, along with the expected profile of the dipole trapping laser (red). The focus of the trapping laser, and position of single atom localisation, is at the central cavity antinode. The trapping light will be imaged after propagating through the cavity mode, to correct for optical aberrations.

However, loading the cavity by an atomic fountain does not scale well with the realisation of cavity-mediated remote atom entanglement. Since this is an inherently stochastic process, the probability that two or more cavities contain a single atom lowers significantly. Therefore, we require a method of localising a single atom within each

cavity. The widely adopted solution for trapping a single neutral atom is by a dipole trap [146]. The physical mechanism for dipole forces on an atom is well explained in literature [51]. However, from an implementation perspective, it uses a highly focused red-detuned laser to establish a conservative trapping profile.

The potential landscape of a dipole trap corresponds directly to the intensity distribution of the driving laser. To create confinement on the length scale of a single atom, it must therefore be tightly focused. As considered in Chapter 4, this established one of the primary restrictions in our cavity design process; by ensuring sufficient optical access to the cavity mode for injection of the trapping light. Previous cavities constructed within our group have lacked this geometry, being loaded by atomic fountain instead. Alternative schemes within optical CQED have used a standing wave dipole trap for intra-cavity atomic localisation [41, 60]. However, in our envisaged atom-cavity system we desire a degree of flexibility that may not be achievable via this method.

Since our dipole trapping method does not require the formation of a standing wave or optical lattice, we have notable freedom to shape laser intensity distribution at its focus. By the use of adaptive optics, specifically a phase spatial light modulator (PSLM), highly bespoke trapping geometries may be established. Demonstrations within our glass cell environment include 4×4 square arrays of trapped atoms, each with a standard occupational lifetime of ≈ 10 seconds. By dynamically updating the PSLM phase pattern these trapping sites may be translated, with single atom transport demonstrated over distances of $\approx 20\mu$m.

In principle, this allows for localising several atoms within the cavity, each at isolated positions within the standing wave cavity mode structure. Alternatively, an atom trap may be translated into the central antinode of the cavity, on detection of its occupation. When combined with the multi-feature resonators discussed in Chapter 6, a set of quasi-independent atom-cavity systems may be created in a highly compact

footprint. However, in our current demonstrations, the dipole trap has been positioned at the centre of the MOT cloud for optimal loading, which is incompatible with a cavity present at this position.

7.5.2 Cold Atom Loading

Since the dipole trap is a conservative potential, it must be loaded with pre-cooled atoms, whose temperature is less than the depth of the trap. As introduced in the previous section, a magneto-optical trap (MOT) is a well-established technique for the cooling and localisation of neutral atoms [145]. From an implementations standpoint, its core components comprise a pair of anti-Helmholtz magnetic field coils and three orthogonal pairs of counter propagating circularly polarised laser beams. However, when a cavity system is installed into the glass cell, the optical access for atomic cooling and trapping is significantly reduced.

To allow for intra-cavity dipole trapping, large apertures were milled into the vertical faces of the cavity module and mounting superstructure, as highlighted in Figure 7.8. This allows the dipole trapping lenses to be located outside of vacuum, for experimental flexibility. However, there is no geometry that allows for the formation of a standard MOT in this region, comprising three pairs of orthogonally propagating beams. Even if the cavity mounting structure was diminished at the expense of vibrational stability, the pyramidal mirror substrates would distort at least one cooling beam. To overcome this, a more compact method of MOT formation must be implemented.

A straightforward and effective method was realised in [60]. This used a planar mirror for MOT beam retro-reflection, allowing all beams to be located on one side of the cavity. A small aperture was located in this mirror for the emanation of atomic flux, which was guided to the cavity mode by harmonic motion in a dipole trapping potential. Once at the position of the cavity mode, a standing wave dipole trap was established to retain atomic position. However, a highly elegant alternative to this approach is

the creation of a pyramidal MOT. This is being developed by a colleague within our group[25] with its natural integration into the cavity superstructure a joint development.

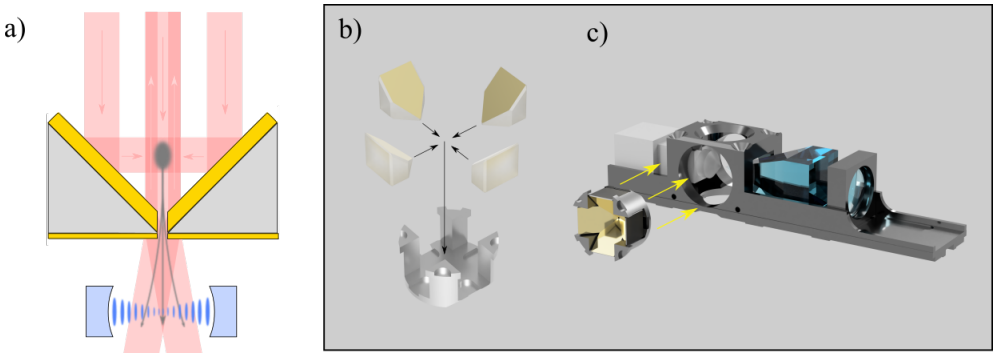

Figure 7.9: **Pyramid MOT. a)** An illustration of the anticipated pyramidal MOT is shown. Gold coated mirrors are illuminated from above, creating the three sets of orthogonal, counter-propagating, red-detuned beams required to establish the cooling action. The atomic cloud is shown in grey. It experiences reduced cooling in the direction of the pyramidal apex. This causes a cold atomic flux to emanate from the pyramidal tip. The pyramid will also be illuminated from the cavity side to control the atomic flux and to form a molasses at the cavity mode, shown in blue. **b-c)** This is installed directly into the cavity module and secured by the superstructure, using an aperture milled into the cavity plate.

A pyramidal MOT comprises four mirrors, whose surfaces are orientated at 45° to a central axis. When the pyramid is illuminated along this axis by a laser beam, three sets of counter-propagating components are directed towards a single axial position. A quadrupole magnetic field is centred on this position, establishing the magneto-optical trap. A small aperture is placed at the apex of the pyramid, which reduces cooling in this direction and creates a stream of cold atoms. A laser beam is directed towards this aperture from the opposite side of the cavity to control the atomic flux and provide cooling at the position of the cavity mode by optical molasses. A schematic of the pyramidal structure and its cooling action is given as Figure 7.9.

At the time of writing, a prototype pyramidal MOT has been fabricated and integrated

[25] Mark IJspeert

into a UHV glass cell environment. Work is ongoing to characterise and optimise the cold atomic flux produced by this structure. It will then be integrated into the dipole trapping glass cell, to establish an updated procedure for atomic loading. Finally, this system will be integrated with the cavity, completing the trapped atom-cavity interface.

7.6 Review and Outlook

In this chapter we performed the design, construction and implementation of a hybrid geometry science cavity, to act as an elementary quantum interface between a trapped neutral atom and single photon. This paired a super-polished near-planar mirror for efficient photon extraction, with an ablated pyramidal micro-mirror to ensure modal confinement. This realised a cavity with strong atom-cavity coupling, $\{g_0, \kappa, \gamma\} = \{11, 2.1, 3\} \times 2\pi$MHz, ample optical access for dipole trapping, $\theta_{oa} > 45°$, photon extraction of $\eta_{ext} \approx 0.5$ and a narrow linewidth of $\Delta\omega_c = 4.2 \times 2\pi$MHz. Our design requirements outlined in Chapter 4 were therefore achieved, with excellent scope to improve cavity properties in future. This would be via the creation a paired pyramidal cavity, once a suitable batch of high transmission mirrors has been fabricated.

A novel process of cavity fabrication was demonstrated, tailored for use with the pyramidal mirrors. This was developed in response to the demanding alignment tolerance of ablated mirrors, without the benefit of integrated beam steering inherent to fibre-tip cavity fabrication. An optomechanical U-Bench structure allowed for a concise alignment process, accurately forming the cavity mode at the anticipated site of atomic trapping. Since this technique ensures that the alignment laser represents the ideal optical axis for the cavity in vacuum, when combined with the integrated mode-matching lens in the cavity module, optical excitation of the cavity and generation of a resonance signal for frequency stabilisation was readily achieved. This will allow the transmissive mirror to be used purely for single photon extraction, with their efficient capture performed using free-space optical elements.

Frequency stabilisation of the cavity to an atom-resonant laser demonstrated a core technical requirement for its use in coherent quantum control. This represents vibrational stability of the pyramidal mirrors, assisted by the cavity mounting superstructure. Future research is expected to include a quantitative analysis of cavity stability and the application of electronic feedback, highlighting the role each mechanical element plays in supporting persistent resonance locking. This will allow direct comparison with alternative cavity geometries and set a benchmark for the fabrication of a paired pyramidal cavity.

The cavity module and mounting structure was specially designed to facilitate a single atom dipole trap, establishing a deterministic atom-cavity interface. A compact optical device required in the formation of a MOT is currently under development. Once this has reached an adequate stage of maturity, it will be integrated with an established platform for re-configurable dipole trapping. Since the cavity, MOT optics and dipole trap have all been installed within identical glass cell environments, their eventual unification should be a straightforward task.

Overall, our science cavity has been designed, fabricated and implemented in a modular and reproducible manner. In the long term, this will assist in the replication of the entire atom-cavity system, for the creation of an elementary quantum network. However, it also permits updating the cavity design, without significant modification to its operational environment. We have highlighted an interest in creating a paired pyramidal mirror cavity, however a viable alternative is the multi-feature resonator demonstrated in Chapter 6. This would require some modification to our alignment technique, to ensure a desired detuning of each respective cavity mode. However, when combined with our re-configurable dipole trap, may lead to a scalable array of atom-cavity interfaces within the same glass cell environment. This would significantly reduce the technical overhead and footprint of the envisaged quantum network, as well as provide a highly

novel platform for performing experimentation in the wider field of optical CQED.

Chapter 8

Conclusion

The objective and outcome of this thesis was to develop new methods in the manufacture of optical cavities for coherent quantum control of trapped neutral atoms and single photons. This has been a principally technical endeavour, requiring increased open access to the cavity mode for atomic localisation, without detriment to its performance as a quantum interface. In this effort, we have described novel contributions to cavity design, ablative glass shaping, optical mirror characterisation, cavity alignment and experimental operation. We shall now provide a brief review of the work of this thesis, before discussing its role in the future direction of our research group and its greater context in the field of optical CQED.

8.1 Review of Progress

Considering single photon production, we outlined the core requirements for cavity suitability and developed metrics of performance by which they could be evaluated. In applying this quantitative analysis to an existing set of cavities, we highlighted a potentially significant performance advantage in migrating to the use of sub-millimetre scale cavity mirrors. In reviewing the state of the art in micro-mirror manufacturing methods, we recommended and elected to pursue laser ablation.

The developed ablation platform combines a number of recent technical developments in micro-patterned ablation, rotational smoothing and integrated surface reconstruction [40, 47]. The current system is exceptionally versatile, able to produce arrays of several hundred ablation features or highly compact mirror geometries suitable for

open access cavity fabrication. We engineered an enhanced control over mirror ellipticity and feature diameter, later demonstrated by a cavity with negligible geometric birefringence and an effective length of $L_{cav} > 1$mm. We presented a new scheme for the laser ablation of pyramidal micro-mirrors, combining the facet geometry of fibre-tips with the experimental practicality of bulk optics.

An apparatus for precise cavity length scanning spectroscopy was developed, allowing for a detailed characterisation of a prototypical batch of pyramidal mirrors. This operated outside of vacuum, permitting a diverse range of spectral measurements to be performed in a time efficient manner. The observed properties of the pyramidal mirrors were consistent with the current state of the art in ablated fibre-tips. The incoherent losses observed, $\mathcal{L}_T = 37$ppm, were consistent with predictions from surface roughness measurements, $\mathcal{L}_S = 29$ppm, and standard coating absorption losses, $\mathcal{L}_A < 10$ppm. Since surface roughness largely derived from our simple method of substrate preparation, mirror quality is anticipated to improve in future work. To assist in its more detailed analysis, a cavity ringdown spectrometer was constructed and demonstrated.

A novel process of cavity fabrication was developed, to ensure a cavity mode is ideally placed for optical interrogation and atomic loading when integrated into vacuum. The optomechanical structure developed for this process greatly simplifies alignment, by realising an equivalent version of the integrated beam steering and cavity mode matching inherent to fibre-tip cavity production. The process is highly repeatable and applicable to a wide range of mirror geometries.

A hybrid pyramidal mirror cavity was fabricated and implemented into vacuum, to act as an elementary interface between light and matter. This was a geometrically asymmetric pairing of mirrors, comprising a $\sigma_m = 100\mu$m diameter pyramidal mirror and high transmission super-polished tapered mirror. Standard requirements for polarised single photon production with ^{87}Rb were met [38, 122], with anticipated

strong atom-cavity coupling, $\{g_0, \kappa, \gamma\} = \{11, 2.1, 3\} \times 2\pi$MHz, narrow cavity linewidth, $\Delta\omega_c = 4.2 \times 2\pi$MHz, and an expected photon extraction efficiency of $\eta_{ext} = 0.5$. Most crucially however, there is excellent open access to the cavity mode to support a single atom dipole trap.

The stabilisation of cavity length to an atom resonant laser demonstrated its future suitability as an operational science cavity, as well as highlighting the inherent vibrational stability of pyramidal mirrors. The integration of atomic cooling and trapping methods into the developed system are ongoing.

8.2 Future Outlook

Ongoing research continues towards a direct demonstration of strong cavity coupling to a single neutral atom in a re-configurable dipole trap. This requires the unification of the fabricated cavity, compact atomic cooling method and dipole trapping optics. Each has been developed in identical vacuum environments and this process is expected to be straightforward.

A strong emphasis on reproducibility and modularity was placed on cavity fabrication. This is expected to support the rapid replication of the trapped atom-cavity system, establishing our intended platform for studying cavity mediated quantum entanglement. This modularity also allows for the cavity design to be updated and refined, without a requirement to redesign the greater experimental architecture. This may be via the production of a paired pyramidal cavity for enhanced atom-cavity coupling strength, or a continuation of the hybrid geometry produced in Chapter 7. Our platform for laser ablation is able to facilitate a broad range of potential mirror substrates, including optical fibres.

Furthermore, with the development of the multi-feature resonators demonstrated in

Chapter 6, a novel route to scaling our atom-cavity quantum network may be realised. In principle, it allows multiple coherent interfaces to be established side-by-side within a single vacuum chamber. A prototype multi-feature cavity may be installed in the short term, using our existing methods of cavity fabrication. However, to engineer control over the respective detuning of each cavity mode, our cavity module design will need to be refined.

This work has focused on the design and implementation of optical cavities for interfacing light and matter. The requirement for open-access cavities with high quantum efficiency is common to a range of qubit platforms. While each species of emitter requires a distinct method of localisation, in the optical domain the demand for significant modal confinement with high quality mirror surfaces is near universal. Indeed, the faithful interconversion of static and travelling qubits constitutes a major challenge in establishing a quantum network with true scalability. The creation of such a network is anticipated to realise a new paradigm in information processing and a significant milestone in the growing development of quantum technologies.

Bibliography

1. Planck, M. Über das Gesetz der Energieverteilung im Normalspectrum. *Annalen der Physik* **4,** 553–563 (1901).
2. Einstein, A. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik* **322,** 132–148 (1905).
3. Lewis, G. N. The Conservation of Photons. *Nature* **188,** 874–875 (1926).
4. Bohr, N. On the Constitution of Atoms and Molecules. *Philos. Mag.* **26,** 1–24 (1913).
5. Heisenberg, W. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Zeitschrift für Physik* **43,** 172–198 (1927).
6. Einstein, A., Podolsky, B. & Rosen, N. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete. *Physical Review* **47,** 2–5 (1935).
7. Ekert, A. K. Quantum cryptography based on Bell's theorem. *Physical Review Letters* **67,** 661–663 (1991).
8. Gisin, N., Ribordy, G., Tittel, W. & Zbinden, H. Quantum cryptography. *Rev. Mod. Phys.* **74,** 145–195 (2002).
9. Feynman, R. P. Simulating Physics with Quantum Computers. *International Journal of Theoretical Physics* **21,** 467–488 (1982).
10. Georgescu, I. M., Ashhab, S. & Nori, F. Quantum simulation. *Rev. Mod. Phys.* **86,** 153–185 (2014).
11. Bouwmeester, D., Ekert, A. K. & Zeilinger, A. *The Physics of Quantum Information* (Springer, Berlin, Heidelberg, 2000).
12. Shor, P. W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing* **26,** 1484–1509 (1997).
13. Grover, L. K. A fast quantum mechanical algorithm for database search. *Proceedings, 28th Annual ACM Symposium on the Theory of Computing (STOC),* 212–219 (1996).
14. DiVincenzo, D. P. The Physical Implementation of Quantum Computation. *Fortschritte der Physik* **48,** 771–783 (2000).
15. Kjaergaard, M. *et al.* Superconducting Qubits: Current State of Play. *Annual Review of Condensed Matter Physics* **11,** 369–395 (2020).
16. Kloeffel, C. & Loss, D. Prospects for Spin-Based Quantum Computing in Quantum Dots. *Annual Review of Condensed Matter Physics* **4,** 51–81 (2013).
17. Nizovtsev, A. P. *et al.* A quantum computer based on NV centers in diamond: Optically detected nutations of single electron and nuclear spins. *Optics and Spectroscopy* **99,** 233–244 (2005).

18. Kok, P. *et al.* Linear optical quantum computing with photonic qubits. *Rev. Mod. Phys.* **79,** 135–174 (2007).
19. Reiserer, A. & Rempe, G. Cavity-based quantum networks with single atoms and optical photons. *Rev. Mod. Phys.* **87,** 1379–1418 (2015).
20. Bruzewicz, C. D., Chiaverini, J., McConnell, R. & Sage, J. M. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews* **6,** 021314 (2019).
21. Kimble, H. J. The quantum internet. *Nature* **453,** 1023–1030 (2008).
22. Kuhn, A. *Cavity Induced Interfacing of Atoms and Light* (Springer, Cham, 2015).
23. Monroe, C. *et al.* Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Phys. Rev. A* **89,** 022317 (2014).
24. Stephenson, L. J. *et al.* High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network. *Phys. Rev. Lett.* **124,** 110501 (2020).
25. Brown, K. R., Kim, J. & Monroe, C. Co-designing a scalable quantum computer with trapped atomic ions. *npj Quantum Information* **2,** 16034 (2016).
26. Kuhn, A., Hennrich, M. & Rempe, G. Deterministic Single-Photon Source for Distributed Quantum Networking. *Phys. Rev. Lett.* **89,** 067901 (2002).
27. Welte, S., Hacker, B., Daiss, S., Ritter, S. & Rempe, G. Photon-Mediated Quantum Gate between Two Neutral Atoms in an Optical Cavity. *Phys. Rev. X* **8,** 011018 (2018).
28. Dilley, J. *A single-photon source for quantum networking* PhD thesis (Oxford University, 2012).
29. Takahashi, H., Kassa, E., Christoforou, C. & Keller, M. Strong Coupling of a Single Ion to an Optical Cavity. *Phys. Rev. Lett.* **124,** 013602 (2020).
30. Podoliak, N., Takahashi, H., Keller, M. & Horak, P. Comparative Numerical Studies of Ion Traps with Integrated Optical Cavities. *Phys. Rev. Applied* **6,** 044008 (2016).
31. Kassa, E., Takahashi, H., Christoforou, C. & Keller, M. Precise positioning of an ion in an integrated Paul trap-cavity system using radiofrequency signals. *Journal of Modern Optics* **65,** 520–528 (2018).
32. Specht, H. P. *et al.* A single-atom quantum memory. *Nature* **473,** 190–193 (2011).
33. Thompson, R. J., Rempe, G. & Kimble, H. J. Observation of normal-mode splitting for an atom in an optical cavity. *Phys. Rev. Lett.* **68,** 1132–1135 (1992).
34. Rempe, G., Thompson, R. J., Kimble, H. J. & Lalezari, R. Measurement of ultralow losses in an optical interferometer. *Opt. Lett.* **17,** 363–365 (1992).
35. Hood, C. J., Kimble, H. J. & Ye, J. Characterization of high-finesse mirrors: Loss, phase shifts, and mode structure in an optical cavity. *Phys. Rev. A* **64,** 033804 (2001).
36. Hunger, D., Deutsch, C., Barbour, R. J., Warburton, R. J. & Reichel, J. Laser micro-fabrication of concave, low-roughness features in silica. *AIP Advances* **2,** 012119 (2012).

37. Dolan, P. R., Hughes, G. M., Grazioso, F., Patton, B. R. & Smith, J. M. Femtoliter tunable optical cavity arrays. *Opt. Lett.* **35,** 3556–3558 (2010).

38. Barrett, T. D., Barter, O., Stuart, D., Yuen, B. & Kuhn, A. Polarization Oscillations in Birefringent Emitter-Cavity Systems. *Phys. Rev. Lett.* **122,** 083602 (2019).

39. Uphoff, M., Brekenfeld, M., Rempe, G. & Ritter, S. Frequency splitting of polarization eigenmodes in microscopic Fabry–Perot cavities. *New Journal of Physics* **17,** 013053 (2015).

40. Takahashi, H. *et al.* Novel laser machining of optical fibers for long cavities with low birefringence. *Opt. Express* **22,** 31317–31328 (2014).

41. Garcia, S., Ferri, F., Ott, K., Reichel, J. & Long, R. Dual-wavelength fiber Fabry-Perot cavities with engineered birefringence. *Opt. Express* **26,** 22249–22263 (2018).

42. Benedikter, J. *et al.* Transverse-mode coupling and diffraction loss in tunable Fabry–Pérot microcavities. *New J. Phys.* **17,** 053051 (2015).

43. Podoliak, N., Takahashi, H., Keller, M. & Horak, P. Harnessing the mode mixing in optical fiber-tip cavities. *J. Phys. B: At. Mol. Opt. Phys.* **50,** 085503 (2017).

44. Hunger, D. *et al.* A fiber Fabry–Perot cavity with high finesse. *New J. Phys.* **12,** 065038 (2010).

45. Gallego, J. *et al.* High-finesse fiber Fabry–Perot cavities: stabilization and mode matching analysis. *Applied Physics B* **122,** 47 (2016).

46. Gulati, G. K., Takahashi, H., Podoliak, N., Horak, P. & Keller, M. Fiber cavities with integrated mode matching optics. *Scientific Reports* **7,** 5556 (2017).

47. Ott, K. *et al.* Millimeter-long fiber Fabry-Perot cavities. *Opt. Express* **24,** 9839–9853 (2016).

48. Janitz, E., Ruf, M., Fontana, Y., Sankey, J. & Childress, L. High mechanical bandwidth fiber-coupled Fabry-Perot cavity. *Opt. Express* **25,** 20932–20943 (2017).

49. Saavedra, C., Pandey, D., Alt, W., Pfeifer, H. & Meschede, D. Tunable fiber Fabry-Perot cavities with high passive stability. *Opt. Express* **29,** 974–982 (2021).

50. Brachmann, J. F. S., Kaupp, H., Hänsch, T. W. & Hunger, D. Photothermal effects in ultra-precisely stabilized tunable microcavities. *Opt. Express* **24,** 21205–21215 (2016).

51. Grimm, R., Weidemüller, M. & Ovchinnikov, Y. B. Optical Dipole Traps for Neutral Atoms. *Advances In Atomic, Molecular, and Optical Physics* **42,** 95–170 (2000).

52. Mabuchi, H., Turchette, Q. A., Chapman, M. S. & Kimble, H. J. Real-time detection of individual atoms falling through a high-finesse optical cavity. *Opt. Lett.* **21,** 1393–1395 (1996).

53. Hood, C. J., Chapman, M. S., Lynn, T. W. & Kimble, H. J. Real-Time Cavity QED with Single Atoms. *Phys. Rev. Lett.* **80,** 4157–4160 (1998).

54. Hennrich, M., Legero, T., Kuhn, A. & Rempe, G. Vacuum-Stimulated Raman Scattering Based on Adiabatic Passage in a High-Finesse Optical Cavity. *Phys. Rev. Lett.* **85,** 4872–4875 (2000).

55. Münstermann, P., Fischer, T., Maunz, P., Pinkse, P. W. H. & Rempe, G. Dynamics of Single-Atom Motion Observed in a High-Finesse Cavity. *Phys. Rev. Lett.* **82,** 3791–3794 (1999).

56. Barrett, T. *Polarised Single-Photons from a Cavity-Enhanced Atom-Light Interface in Photonic Quantum Networks* PhD thesis (Oxford University, 2018).

57. Ye, J., Vernooy, D. W. & Kimble, H. J. Trapping of Single Atoms in Cavity QED. *Phys. Rev. Lett.* **83,** 4987–4990 (1999).

58. Reiserer, A., Nölleke, C., Ritter, S. & Rempe, G. Ground-State Cooling of a Single Atom at the Center of an Optical Cavity. *Phys. Rev. Lett.* **110,** 223003 (2013).

59. Gallego, J. *et al.* Strong Purcell Effect on a Neutral Atom Trapped in an Open Fiber Cavity. *Phys. Rev. Lett.* **121,** 173603 (2018).

60. Nußmann, S. *et al.* Submicron Positioning of Single Atoms in a Microcavity. *Phys. Rev. Lett.* **95,** 173602 (2005).

61. Kuhr, S. *et al.* Deterministic Delivery of a Single Atom. *Science* **293,** 278–280 (2001).

62. Garcia, S., Ferri, F., Reichel, J. & Long, R. Overlapping two standing waves in a microcavity for a multi-atom photon interface. *Opt. Express* **28,** 15515–15528 (2020).

63. Thompson, J. D. *et al.* Coupling a Single Trapped Atom to a Nanoscale Optical Cavity. *Science* **340,** 1202–1205 (2013).

64. Samutpraphoot, P. *et al.* Strong Coupling of Two Individually Controlled Atoms via a Nanophotonic Cavity. *Phys. Rev. Lett.* **124,** 063602 (2020).

65. Stuart, D. & Kuhn, A. Single-atom trapping and transport in DMD-controlled optical tweezers. *New Journal of Physics* **20,** 023013 (2018).

66. Holland, N. *Holographically Generated Optical Tweezers for the Control of Single Atoms* PhD thesis (Oxford University, 2019).

67. Einstein, A. Zur Quantentheorie der Strahlung. *Phys. Z.* **18,** 121–128 (1917).

68. Haroche, S. & Raimond, J.-M. *Exploring the quantum* (Oxford University Press, 2019).

69. Purcell, E. M., Torrey, H. C. & Pound, R. V. Resonance Absorption by Nuclear Magnetic Moments in a Solid. *Phys. Rev.* **69,** 37–38 (1946).

70. Kleppner, D. Inhibited Spontaneous Emission. *Phys. Rev. Lett.* **47,** 233–236 (1981).

71. Jaynes, E. T. & Cummings, F. W. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE* **51,** 89–109 (1963).

72. Agarwal, G. S. Vacuum-field Rabi oscillations of atoms in a cavity. *J. Opt. Soc. Am. B* **2,** 480–485 (1985).

73. Siegman, A. E. *Lasers* (Univ. Science Books, 1986).

74. Brooker, G. *Modern classical optics* (Oxford University Press, 2014).

75. Grynberg, G., Aspect, A., Fabre, C. & Cohen-Tannoudji, C. *Introduction to Quantum Optics: From the Semi-classical Approach to Quantized Light* (Cambridge University Press, 2010).

76. Fox, M. *Quantum optics* (Oxford University Press, 2013).

77. Perot, A. & Fabry, C. On the Application of Interference Phenomena to the Solution of Various Problems of Spectroscopy and Metrology. *The Astrophysical Journal* **9,** 87 (1899).

78. Ismail, N., Kores, C. C., Geskus, D. & Pollnau, M. Fabry-Perot resonator: spectral line shapes, generic and related Airy distributions, linewidths, finesses, and performance at low or frequency-dependent reflectivity. *Opt. Express* **24,** 16366–16389 (2016).

79. Kimble, H. J. Strong Interactions of Single Atoms and Photons in Cavity QED. *Physica Scripta* **T76,** 127 (1998).

80. Nisbet-Jones, P. B. R., Dilley, J., Holleczek, A., Barter, O. & Kuhn, A. Photonic qubits, qutrits and ququads accurately prepared and delivered on demand. *New Journal of Physics* **15,** 053007 (2013).

81. Agarwal, G. S. *Quantum Optics* (Cambridge University Press, 2012).

82. Gibbs, H. *Optical Bistability: Controlling Light With Light* (Elsevier Science, 2012).

83. Boca, A. *et al.* Observation of the Vacuum Rabi Spectrum for One Trapped Atom. *Phys. Rev. Lett.* **93,** 233603 (2004).

84. Bergmann, K. *et al.* Roadmap on STIRAP applications. *J. Phys. B: At. Mol. Opt. Phys.* **52,** 202001 (2019).

85. Dilley, J., Nisbet-Jones, P., Shore, B. W. & Kuhn, A. Single-photon absorption in coupled atom-cavity systems. *Phys. Rev. A* **85,** 023834 (2012).

86. McKeever, J. Deterministic Generation of Single Photons from One Atom Trapped in a Cavity. *Science* **303,** 1992–1994 (2004).

87. Vasilev, G. S., Ljunggren, D. & Kuhn, A. Single photons made-to-measure. *New Journal of Physics* **12,** 063024 (2010).

88. Holleczek, A. *et al.* Quantum Logic with Cavity Photons From Single Atoms. *Phys. Rev. Lett.* **117,** 023602 (2016).

89. Barrett, T. D. *et al.* Multimode interferometry for entangling atoms in quantum networks. *Quantum Science and Technology* **4,** 025008 (2019).

90. Wilk, T., Webster, S. C., Specht, H. P., Rempe, G. & Kuhn, A. Polarization-Controlled Single Photons. *Phys. Rev. Lett.* **98,** 063601 (2007).

91. Wilk, T., Webster, S. C., Kuhn, A. & Rempe, G. Single-Atom Single-Photon Quantum Interface. *Science* **317,** 488–490 (2007).

92. Ritter, S. *et al.* An elementary quantum network of single atoms in optical cavities. *Nature* **484,** 195–200 (2012).

93. Kogelnik, H. & Li, T. Laser Beams and Resonators. *Appl. Opt.* **5,** 1550–1567 (1966).

94. Silfvast, W. T. *Laser Fundamentals* (Cambridge University Press, 2008).

95. Hercher, M. The Spherical Mirror Fabry-Perot Interferometer. *Applied Optics* **7,** 951 (1968).

96. Durak, K., Nguyen, C. H., Leong, V., Straupe, S. & Kurtsiefer, C. Diffraction-limited Fabry–Perot cavity in the near concentric regime. *New Journal of Physics* **16,** 103002 (2014).

97. Laures, P. Geometrical Approach to Gaussian Beam Propagation. *Appl. Opt.* **6,** 747–755 (1967).

98. Bandres, M. A. & Gutiérrez-Vega, J. C. Ince–Gaussian modes of the paraxial wave equation and stable resonators. *J. Opt. Soc. Am. A* **21,** 873–880 (2004).

99. Bandres, M. A. & Gutiérrez-Vega, J. C. Ince–Gaussian beams. *Opt. Lett.* **29,** 144–146 (2004).

100. Schwarz, U. T., Bandres, M. A. & Gutiérrez-Vega, J. C. Observation of Ince–Gaussian modes in stable resonators. *Opt. Lett.* **29,** 1870–1872 (2004).

101. Feng, S. & Winful, H. G. Physical origin of the Gouy phase shift. *Opt. Lett.* **26,** 485–487 (2001).

102. Macleod, H. A. *Thin-film optical filters* (Taylor & Francis, 1986).

103. Sheppard, C. J. R. Approximate calculation of the reflection coefficient from a stratified medium. *Pure and Applied Optics: Journal of the European Optical Society Part A* **4,** 665–669 (1995).

104. Gangloff, D. *et al.* Preventing and reversing vacuum-induced optical losses in high-finesse tantalum (V) oxide mirror coatings. *Opt. Express* **23,** 18014–18028 (2015).

105. Bennett, J. M. Recent developments in surface roughness characterization. *Measurement Science and Technology* **3,** 1119–1127 (1992).

106. Chao, S., Wang, W.-H. & Lee, C.-C. Low-loss dielectric mirror with ion-beam-sputtered TiO_2-SiO_2 mixed films. *Appl. Opt.* **40,** 2177–2182 (2001).

107. Apfel, J. H. Optical coating design with reduced electric field intensity. *Appl. Opt.* **16,** 1880–1885 (1977).

108. Saharyan, A., Alvarez, J., Doherty, T. H., Kuhn, A. & Guerin, S. Light-matter interaction in open cavities with dielectric stacks. *Appl. Phys. Lett.* **118,** 154002 (2021).

109. Zoller, P. *et al.* Quantum information processing and communication. *The European Physical Journal D* **36,** 203–228 (2005).

110. Trichet, A. A. P., Dolan, P. R., Coles, D. M., Hughes, G. M. & Smith, J. M. Topographic control of open-access microcavities at the nanometer scale. *Opt. Express* **23,** 17205–17216 (2015).

111. Reyntjens, S. & Puers, R. A review of focused ion beam applications in microsystem technology. *Journal of Micromechanics and Microengineering* **11,** 287–300 (2001).

112. Nemoto, K. *et al.* Photonic Quantum Networks formed from NV- centers. *Scientific Reports* **6,** 26284 (2016).

113. Petrak, B. *et al.* Isotopic gas analysis through Purcell cavity enhanced Raman scattering. *Applied Physics Letters* **108,** 091107 (2016).

114. Lather, J., Bhatt, P., Thomas, A., Ebbesen, T. W. & George, J. Cavity Catalysis by Cooperative Vibrational Strong Coupling of Reactant and Solvent Molecules. *Angewandte Chemie International Edition* **58,** 10635–10638 (2019).

115. Ruelle, T., Poggio, M. & Braakman, F. Optimized single-shot laser ablation of concave mirror templates on optical fibers. *Appl. Opt.* **58,** 3784–3789 (2019).

116. Dong, J. *Laser trapping of atoms and cavity quantum electrodynamics in fibre-tip microcavities* PhD thesis (Oxford University, 2014).

117. Walker, B. T. *et al.* Driven-dissipative non-equilibrium Bose–Einstein condensation of less than ten photons. *Nature Physics* **14,** 1173–1177 (2018).

118. Nguyen, C. H. *et al.* Single atoms coupled to a near-concentric cavity. *Phys. Rev. A* **96,** 031802 (2017).

119. Harlander, M., Brownnutt, M., Hänsel, W. & Blatt, R. Trapped-ion probing of light-induced charging effects on dielectrics. *New Journal of Physics* **12,** 093035 (2010).

120. Barter, O. *Deterministic Quantum Feedback Control in Probabilistic Atom-Photon Entanglement* PhD thesis (Oxford University, 2016).

121. Steck, D. A. Rubidium 87 D Line Data. Available online at: `https://steck.us/alkalidata/rubidium87numbers.pdf` (Nov. 2019).

122. Barrett, T. D., Stuart, D., Barter, O. & Kuhn, A. Nonlinear Zeeman effects in the cavity-enhanced emission of polarised photons. *New Journal of Physics* **20,** 073030 (2018).

123. Nisbet-Jones, P. *Shaping single photons* PhD thesis (Oxford University, 2012).

124. Gallego, J. C. *Strong Coupling between Small Atomic Ensembles and an Open Fiber Cavity* PhD thesis (University of Bonn, 2017).

125. Barrett, T. D., Doherty, T. H. & Kuhn, A. Pushing Purcell enhancement beyond its limits. *New Journal of Physics* **22,** 063013 (2020).

126. Bielsa, F. *et al.* Birefringence of interferential mirrors at normal incidence. *Applied Physics B* **97,** 457–463 (2009).

127. Kassa, E., Hughes, W., Gao, S. & Goodwin, J. F. Effects of cavity birefringence on remote entanglement generation. Preprint Available: arXiv:2008.11712 [quant-ph] (2020).

128. Chichkov, B. N., Momma, C., Nolte, S., Alvensleben, F. & Tünnermann, A. Femtosecond, picosecond and nanosecond laser ablation of solids. *Applied Physics A Materials Science & Processing* **63,** 109–115 (1996).

129. Stuart, B. C. *et al.* Nanosecond-to-femtosecond laser-induced breakdown in dielectrics. *Phys. Rev. B* **53,** 1749–1761 (1996).

130. Mirza, I. *et al.* Ultrashort pulse laser ablation of dielectrics: Thresholds, mechanisms, role of breakdown. *Scientific Reports* **6,** 39133 (2016).

131. Cui, J.-M. *et al.* Polarization nondegenerate fiber Fabry-Perot cavities with large tunable splittings. *Applied Physics Letters* **112,** 171105 (2018).

132. Nowak, K. M., Baker, H. J. & Hall, D. R. Analytical model for CO2 laser ablation of fused quartz. *Appl. Opt.* **54,** 8653–8663 (2015).

133. Gomez Velez, J. & Muller, A. Purcell-enhanced microcavity Raman scattering from pressurized gases. *Applied Physics Letters* **112,** 041107 (2018).

134. Niklas, C., Wackerbarth, H. & Ctistis, G. A Short Review of Cavity-Enhanced Raman Spectroscopy for Gas Analysis. *Sensors* **21,** 1698 (2021).

135. De Groot, P. Principles of interference microscopy for the measurement of surface topography. *Adv. Opt. Photon.* **7,** 1–65 (2015).

136. Eaton, P. J. & West, P. *Atomic force microscopy* (Oxford University Press, 2019).

137. Puppe, T., Maunz, P., Fischer, T., Pinkse, P. & Rempe, G. Single-Atom Trajectories in Higher-Order Transverse Modes of a High-Finesse Optical Cavity. *Physica Scripta* **T112,** 7 (2004).

138. Brekenfeld, M., Niemietz, D., Christesen, J. D. & Rempe, G. A quantum network node with crossed optical fibre cavities. *Nature Physics* **16,** 647–651 (2020).

139. Degallaix, J. OSCAR a Matlab based optical FFT code. *Journal of Physics: Conference Series* **228,** 012021 (2010).

140. Stuart, D. *Manipulating single atoms with optical tweezers* PhD thesis (Oxford University, 2014).

141. Kawasaki, A. *et al.* Geometrically asymmetric optical cavity for strong atom-photon coupling. *Phys. Rev. A* **99,** 013437 (2019).

142. Mohammed, M. *A Fibre-Tip Fabry-Perot Cavity for Deterministic Strong Atom-Photon Interactions* Unpublished. PhD thesis (Oxford University, 2018).

143. Drever, R. W. P. *et al.* Laser phase and frequency stabilization using an optical resonator. *Applied Physics B Photophysics and Laser Chemistry* **31,** 97–105 (1983).

144. Black, E. D. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics* **69,** 79–87 (2001).

145. Foot, C. J. *Atomic Physics* (Oxford University Press, 2005).

146. Schlosser, N., Reymond, G., Protsenko, I. & Grangier, P. Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature* **411,** 1024–1027 (2001).