Rheinische Friedrich-Wilhelms-Universität
Bonn

Coupled high-finesse optical Fabry-Pérot microcavities

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von
Steffen Wilhelm Gohlke
aus
Siegburg, Deutschland

Bonn, Januar 2026

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der Rheinischen Friedrich-Wilhelms-Universität Bonn.

Gutachter/Betreuer:	Prof. Dr. Michael Köhl
Gutachter:	Prof. Dr. Daqing Wang
Tag der Promotion:	21.05.2026
Erscheinungsjahr:	2026

Abstract

Dissipation is unavoidable in any realistic physical system and can strongly reshape its dynamics and spectra. Coupled cavity systems offer a minimal and highly controllable setting in which coherent photon exchange competes with radiative loss, enabling the study of interference-based control of the system's optical response.

In this thesis, we report the experimental realisation and characterisation of such an open photonic system, consisting of two coherently coupled Fabry–Pérot microcavities. The system is built from a single fibre-based resonator that is split by a high-reflectivity membrane mirror inserted at the cavity centre. Transmission through the membrane enables coherent photon exchange between the sub-cavities. This coherent coupling competes with radiative losses set by the finite mirror reflectivity, resulting in hybridised modes that extend across both resonators. In this way, the system forms a photonic molecule that provides an all-optical analogue of linear light–matter interaction, offering a controlled setting to study how coherent coupling and dissipation shape the optical response of an open system.

A central contribution of this work is the development of a dielectric-coated silicon-nitride membrane mirror that combines high reflectivity (transmittance $\mathcal{T} = 452(47)$ ppm) with sufficient surface quality to sustain high-finesse resonances on both sides. Near degeneracy of these resonances, we resolve a pronounced avoided crossing, demonstrating strong coherent coupling $g = 2\pi \times 2.859(2)$ GHz. Beyond this regime, we achieve wide tunability of the effective coupling by exploiting higher-order transverse cavity modes, reducing the coupling strength by nearly two orders of magnitude, which allows us to examine the transition from strongly hybridised supermodes to weakly perturbed single-cavity resonances. Using dynamic cavity-length modulation to emulate an imbalance of loss rates, we additionally observe strong interference effects in the steady-state optical response, manifesting as EIT-like narrow transparency features.

We further identify absorption in the silicon-nitride membrane that induces a nonlinear response in one sub-cavity and gives rise to optical bistability. Based on these measurements, we discuss how the photonic molecule picture extends into the nonlinear regime, where interference features become photon-number dependent. We outline conditions under which nonlinear interference could suppress multiphoton excitation via the unconventional photon blockade mechanism, and discuss a realistic path toward this effect by leveraging transverse access to one cavity site for integrating additional nonlinear media.

Parts of this thesis have been published or will be published in the following articles:

Steffen Gohlke et al. “Coupled high-finesse optical Fabry-Perot microcavities”. *Phys. Rev. A* **109**, L011501 (2024). DOI: 10.1103/PhysRevA.109.L011501

Contents

1. Introduction **1**

2. Fundamentals of optical cavity resonators **5**
2.1. The Fabry–Pérot etalon 6
2.1.1. Limit of high mirror reflectivities 10
2.2. Transversal electromagnetic (TEM) cavity modes 11
2.2.1. Fundamental Gaussian beam 11
2.2.2. Boundary condition of a resonator with curved end mirrors 12
2.2.3. Higher order transversal electromagnetic wave solutions of an optical cavity resonator 15
2.3. Experimental realisation of a fibre-based cavity apparatus 17
2.3.1. Measurement of the reflection spectrum of the fibre cavity 18
2.3.2. Limitations due to transverse cavity mode effects 20
2.3.3. Summary 23

3. Production and characterisation of a dielectric membrane mirror **25**
3.1. Geometry of the membrane mirror 26
3.1.1. Interferometry of the membrane surface 27
3.1.2. Modifications to the membrane geometry 28
3.2. Optimisation of the mirror coating composition 33
3.2.1. Refractive index of silicon nitride 33
3.2.2. Optimised coating composition 36
3.3. Optical test cavity with a membrane mirror 37
3.3.1. Absorption within the silicon nitride layer 40
3.4. Investigation of the membrane's vibrational modes 42
3.4.1. Detection of the membrane motion 43
3.4.2. Mechanical model 46
3.5. Summary 50

4. Realisation of two optical cavity resonators coupled via shared end mirror **51**
4.1. Mathematical description of two coupled resonators 52
4.1.1. Mode splitting 55
4.1.2. Estimation of the coupling strength 56
4.2. Experimental implementation of the mirror-in-the-middle setup 57
4.2.1. Final sub-cavity performance 62

4.3. Spectroscopy and characterisation near degeneracy point 63
4.3.1. Scheme of measurement 64
4.3.2. Measurement of the avoided crossing 68
4.4. Summary and future improvements 70

5. Tuneable coupling strength by higher-order mode mixing 71
5.1. Mode overlap of higher-order transversal cavity modes 73
5.2. A purely optical analogue to the Purcell effect 75
5.2.1. Dissipative interaction in the intermediate and strong coupling regime 78
5.3. Formation of the minimal mode splitting 81
5.4. Range of accessible coupling strengths 83
5.5. Simulation of coupled-resonator-induced-transparencies 85
5.6. Summary 88

6. Absorption-induced nonlinearities in a coupled-cavity environment 89
6.1. Characterisation of the nonlinear behaviour in a single cavity resonator . 90
6.1.1. Mathematical model 91
6.1.2. Measurements of the bistable optical response 93
6.1.3. Response time of the absorption-induced nonlinearity 98
6.2. Towards unconventional photon blockade effect in coupled cavities 103
6.2.1. Optical response of the nonlinear photonic molecule 104
6.2.2. Optimal parameters for destructive interference 108
6.2.3. Prospects for experimental implementation 111
6.3. Summary 112

7. Summary and Future Direction 115
7.1. Outlook: Temporally resolved photon tunnelling between cavity sites . . . 116

A. Introduction to the transfer-matrix method 119
A.1. Reflection and transmission spectra of the multilayer stack 119
A.2. Intensity distribution within the mirror coating 123

Bibliography 127

Acknowledgements 135

1. Introduction

Quantum systems are often modelled as perfectly isolated. In reality, interaction with the environment is unavoidable and leads to effectively non-unitary dynamics and mixed quantum states [1–3]. This interaction causes information to leak into the environment, rendering the system an open quantum system. At the same time, any drive or readout necessarily relies on this very environmental coupling. Therefore, understanding and controlling open-system dynamics is not only foundational but also technologically crucial. One of the central questions here is how dissipation, the irreversible exchange of energy with the environment, competes with coherent evolution, thereby reshaping both states and spectra. This line of questioning has led from early studies of dissipative effects toward modern "open-system engineering", which asks whether tailored dissipation can be exploited to stabilise and even enhance coherence [4, 5].

Optical resonators provide a particularly useful setting to study open-system dynamics. A cavity confines a small set of electromagnetic modes while coupling them to the environment through well-defined ports. This makes the interplay of coherent evolution and photon loss experimentally accessible and quantitatively characterisable via measurements of the cavity's optical response [6, 7]. Such confinement can be realised in ring resonators, where light circulates in a looped waveguide and couples to external channels through evanescent interaction with adjacent waveguides [8]. Alternatively, mirror-based Fabry–Pérot resonators confine light between reflective surfaces, where dissipation is introduced by the finite mirror reflectivity. Fibre-based microcavities constitute a compact realisation of such Fabry–Pérot resonators, formed by machined and coated end facets of two opposing single-mode fibres [9, 10]. In this way, fibre cavities retain transverse access to the optical mode as in macroscopic mirror-based cavities, while offering the waveguide-compatible input and readout of ring resonators.

While a single cavity already provides a controlled balance of confinement and loss, coupling two resonators introduces an additional coherent process, namely photon exchange between the modes. This competition between coherent exchange and leakage directly shapes the optical response of coupled-cavity systems, thereby providing a natural platform for studying open-system dynamics [11]. For weak coupling, the resonances remain close to those of the individual cavities and are only slightly perturbed. As the inter-cavity coupling increases, the relevant excitations are no longer confined to a single resonator. Instead, the system supports delocalised supermodes, meaning hybridised field distributions that extend across both resonators, which appear as distinct features in the optical response. Through this hybridisation, the fields emitted by the two resonators can interfere, creating bright and dark supermodes, enabling engineered dissipation and characteristic features such as narrow transparency windows and asym-

metric lineshapes [12–14]. In this sense, coupled cavities form a photonic molecule that realises a highly tuneable, all-optical analogue of linear light–matter interaction [15–17].

Such photonic molecules have been realised in a range of platforms, including coupled ring resonators [18, 19] and macroscopic mirror-based cavity systems [20]. A coupled-cavity realisation based on fibre-based microcavities was previously demonstrated by inserting a silicon nitride membrane into the cavity centre [21], but the membrane reflectivity in this approach was not sufficient to sustain two well-confined optical modes on either side of the coupling element. Recent advances in high-reflectivity membrane mirrors, achieved for example through patterned nanostructures [20, 22, 23] or dielectric mirror coatings [24, 25], provide a route to overcome this limitation. Combining such membranes with fibre microcavities therefore enables a fully fibre-integrated platform for coupled-cavity experiments and dissipation engineering.

This fibre-integrated cavity geometry is also particularly attractive for nonlinear extensions. Transverse access to the optical mode enables the integration of a wide range of active or nonlinear media, including single atoms or ions, quantum gases, liquid media, and optomechanical elements [26–30], while the small mode volume and strong field enhancement of fibre microcavities can substantially increase the effective interaction strength [10]. Such cavity-enhanced nonlinearities can render the optical response sensitive to the intracavity intensity and, in the quantum regime, to the photon number, providing routes towards all-optical switching and few-photon nonlinear optics, ultimately enabling the generation of nonclassical light [31].

Coupling such a nonlinear cavity to a second optical mode forms a nonlinear photonic molecule, in which nonlinear response, interference, and dissipation can be studied in a controlled manner [32, 33]. In this setting, nonlinear shifts can make the interference between the emitted fields photon-number dependent, so that transparency features of the linear system discriminate between single- and multiphoton states. When tuned appropriately, this mechanism can yield antibunched light even for comparatively weak nonlinearities, effectively suppressing multiphoton transmission while retaining a strong single-photon contribution [34–36]. In this way, interference amplifies the impact of the nonlinearity, providing a more efficient path to nonclassical emission than approaches that rely solely on large intrinsic nonlinearities [37]. It therefore positions fibre-integrated coupled-cavity platforms as especially appealing candidates for compact sources of nonclassical light that are directly compatible with photonic network architectures [38].

In this thesis, we experimentally realise a fibre-based Fabry–Pérot microcavity platform built from fibre mirrors that are fabricated in-house, with their end facets machined within our group. To implement coupling, we develop a coated silicon nitride membrane mirror with sufficiently high reflectivity to split the resonator into two sub-cavities that are coupled via transmission through the membrane. The resulting platform is characterised by demonstrating controlled tuning of the effective coupling strength, revealing signatures ranging from weak perturbations of the uncoupled resonances to the formation of delocalised normal modes. Furthermore, dynamic cavity modulations are shown to manifest as narrow transparency features in the reflected optical response. Finally, nonlinear behaviour arising from absorption in the silicon nitride layer is discussed.

Thesis structure

The objective of this work is to present the experimental realisation and comprehensive characterisation of two coupled high-finesse optical cavities forming a photonic molecule. The platform is based on a fibre Fabry–Pérot microcavity that is divided by a high-reflectivity membrane mirror into two radiatively coupled sub-cavities.

Chapter 2 introduces the fundamentals of optical cavity resonators and establishes the parameters and concepts used throughout the thesis. Particular emphasis is also placed on transverse cavity modes and their experimental relevance. Within this framework, fibre-based cavities are discussed, and limitations associated with transverse-mode coupling are analysed in detail.

Chapter 3 focuses on the fabrication and characterisation of the membrane mirror that transforms the "membrane-in-the-middle" configuration into the "mirror-in-the-middle" regime. The chapter describes the underlying design considerations and technical implementation that lead to the membrane mirror prototype. In addition, single-cavity measurements incorporating the membrane are presented to assess its optical performance and to characterise its mechanical properties.

Chapter 4 describes the integration of the membrane mirror into a pre-aligned fibre cavity and the measurement approach used to quantify coupling between the resulting sub-cavities. The formation of coupled normal modes is identified through the observation of an avoided crossing, providing evidence for coherent photon exchange.

Chapter 5 extends this analysis by tuning the inter-cavity coupling in discrete steps, spanning from weak- to strong-coupling, utilising a lifted orthonormality of the cavity modes, caused by an intrinsic misalignment of the cavity. The resulting spectra are discussed in the photonic-molecule picture relating them to phenomena familiar from matter–light settings. Furthermore, dynamic modulation of the individual cavity lengths is used to realise induced-transparency-like features in a purely optical two-mode system.

Chapter 6 examines nonlinear effects arising from absorption in the silicon nitride membrane that are identified during the membrane characterisation. The thermally induced nonlinearity is first discussed in the single-cavity setting and then generalised to the coupled-cavity configuration to evaluate its influence on the system response. On this basis, the chapter discusses modifications that could strengthen and control the effective nonlinearity, with a prospect to future investigations of photon–photon interactions in this architecture.

Finally, Chapter 7 provides an outlook on potential extensions of the platform, facilitating direct observation of photon tunnelling dynamics in real time, which would offer an additional route to probing the time-domain behaviour of open coupled resonator systems [39].

2. Fundamentals of optical cavity resonators

Contents

2.1. The Fabry–Pérot etalon 6

2.1.1. Limit of high mirror reflectivities 10

2.2. Transversal electromagnetic (TEM) cavity modes 11

2.2.1. Fundamental Gaussian beam 11

2.2.2. Boundary condition of a resonator with curved end mirrors . . 12

2.2.3. Higher order transversal electromagnetic wave solutions of an optical cavity resonator 15

2.3. Experimental realisation of a fibre-based cavity apparatus . 17

2.3.1. Measurement of the reflection spectrum of the fibre cavity . . 18

2.3.2. Limitations due to transverse cavity mode effects 20

2.3.3. Summary 23

Optical cavities are fundamental building blocks in photonics and quantum optics, providing controlled confinement of light in a well-defined mode structure. In a mirror-based Fabry–Pérot geometry, confinement is achieved by two highly reflective mirrors, where the resonance condition is set by the mirror separation (cavity length) and the optical wavelength. Finite mirror reflectivity and residual scattering or absorption introduce photon loss and thus naturally realise an engineered dissipative system [6, 7]. Since the coupled-cavity photonic molecule studied in this thesis is built from such resonators, a detailed understanding of a single cavity provides a solid foundation for the analysis in the following chapters.

Therefore, this chapter establishes the theoretical framework required to analyse coupled Fabry–Pérot cavities. We derive the optical response of a single cavity, i.e. its reflection and transmission spectra, and introduce the key parameters used throughout the thesis. Beyond longitudinal resonances, we also investigate transverse modes, which correspond to variations of the field distribution perpendicular to the optical axis. Including transverse modes yields a more accurate description of real cavities, since they impose additional boundary conditions on the mirror surfaces that govern resonator stability and constrain the viable combinations of cavity length and mirror curvature. Moreover, considering higher-order transverse modes will later prove essential for tuning cavity–cavity coupling, as discussed in Chapter 5.

For the experimental implementation, we employ an in-house fabricated fibre-based microcavity, a Fabry–Pérot resonator that is formed between the end facets of two opposing single-mode fibres coated with high-reflectivity dielectric mirrors [9, 10]. This microcavity serves as the starting point that will later be split by a membrane mirror into two coupled sub-cavities. We also characterise this fibre-based microcavity setup as a baseline for the platform extension (presented in Chapter 4) by examining its longitudinal and transverse mode structure. Based on these results, we discuss intrinsic limitations of the fibre-cavity architecture [40]. We then use this analysis to extract reference parameters to guide subsequent design choices, particularly with regard to the geometric requirements of the membrane mirror needed to realise the coupled-cavity configuration.

2.1. The Fabry–Pérot etalon

The Fabry–Pérot cavity is one of the most fundamental and widely employed optical cavity resonators. It consists of two reflective mirrors positioned opposite each other at a fixed separation. Due to the finite reflectivity of these mirrors, an incident light field partially transmits into the cavity, where it undergoes multiple reflections between the mirrors. When the separation between the mirrors ensures that the optical phase accumulated over a single round trip corresponds to an integer multiple of 2π, constructive interference occurs. This results in the formation of a standing-wave pattern and a resonant enhancement of the oscillating intracavity field.

A schematic representation of a Fabry–Pérot cavity resonator is provided in Figure 2.1.(a), illustrating a cavity formed between two parallel mirrors separated by a distance d. The trajectory of a light beam incident from the left is also depicted. Upon transmission through the left mirror, the light field circulates within the cavity, creating an overlap between forward-propagating field components and their counter-propagating reflections. To develop an analytical framework for describing this interference effect, it is useful to express the field as a superposition of these two counter-propagating components. Here, we denote the forward-propagating component of the circulating field as a_{circ}, while the backward-propagating component is labeled as a'_{circ}, collectively forming the vector $\boldsymbol{a}_{\mathrm{circ}} = (a_{\mathrm{circ}}, a'_{\mathrm{circ}})^{\mathrm{T}}$. Throughout this derivation, optical fields at different positions along the cavity axis generally follow the same notation.

In general, any optical element can be mathematically represented by a transformation matrix that relates those electric-field vectors on either side of the element. More complex optical systems can thus be modelled as a sequence of such transformations[1]. Consequently, the Fabry–Pérot cavity can be decomposed into a series of optical interactions: an initial reflection and transmission at the first mirror, propagation over a distance d to the second mirror, and a concluding interaction at the second mirror. This decomposition is illustrated in Figure 2.1.(b), which highlights the sequential interactions governing the optical response of the cavity.

[1] In the literature (e.g., Ref. [41]), this approach is commonly referred to as the *transfer-matrix method*. A more general formalism is outlined in Appendix A.

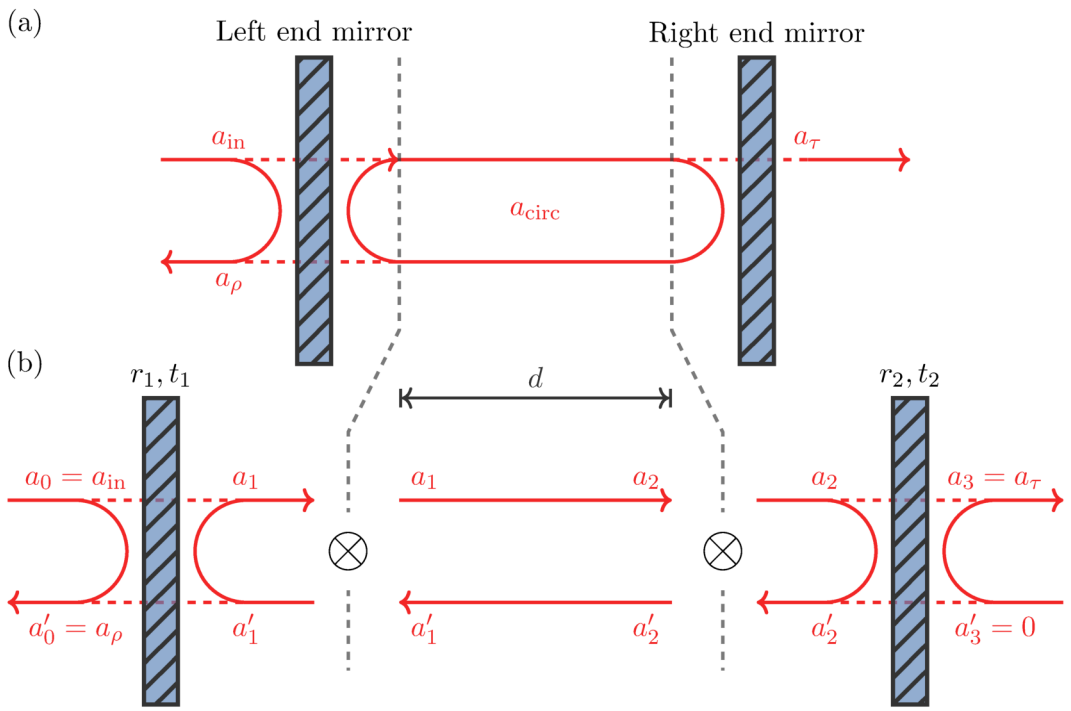


Figure 2.1.: (a) Schematic of a Fabry–Pérot cavity formed by two planar mirrors separated by a distance d. The circulating optical path is indicated in red, illustrating the self-interference of the intracavity field. (b) Decomposition of the cavity into its individual optical components acting on an incident field vector $\boldsymbol{a}_{\text{in}}$. The vector entries denote field components propagating in the positive and negative directions along the cavity axis and determine the interference of the counter-propagating fields.

To characterise the effect of the first mirror on the incident field vector $\boldsymbol{a}_0$, we introduce the transformation matrix $\mathcal{M}_1$, which relates the field components on either side of the mirror, as illustrated in Figure 2.1.(b). This transformation is given by:

$$\begin{pmatrix} a_1 \\ a'_1 \end{pmatrix} = \underbrace{\frac{1}{t_1}\begin{pmatrix} r_1^2 + t_1^2 & ir_1 \\ -ir_1 & 1 \end{pmatrix}}_{\mathcal{M}_1} \cdot \begin{pmatrix} a_0 \\ a'_0 \end{pmatrix}, \tag{2.1}$$

where $\boldsymbol{a}_1 = (a_1, a'_1)^{\mathrm{T}}$ represents the field vector on the opposite side of the mirror, and r_1 and t_1 denote the reflection and transmission amplitude coefficients of the mirror, respectively. The imaginary unit i accounts for the relative phase between reflected and transmitted field components, resulting in $\arg(r) = \arg(t)$ [41].

Following this initial mirror interaction, the propagation of the fields between the two mirrors is governed by free phase evolution. This propagation is described by the

transformation matrix $\mathcal{P}$:

$$\begin{pmatrix} a_2 \\ a_2' \end{pmatrix} = \underbrace{\begin{pmatrix} \mathrm{e}^{-ikd} & 0 \\ 0 & \mathrm{e}^{ikd} \end{pmatrix}}_{\mathcal{P}} \cdot \begin{pmatrix} a_1 \\ a_1' \end{pmatrix}, \tag{2.2}$$

where k is the magnitude of the wavevector of the incident field, defined as $k = 2\pi n/\lambda_0$, with λ_0 being the vacuum wavelength and n the refractive index of the medium inside the cavity. The phase terms differ in sign due to the opposite propagation directions of the counter-propagating fields.

Completing the sequence with another mirror interaction, represented by $\mathcal{M}_2$, the total effect of the Fabry–Pérot cavity on the incident field vector is described by the composite transformation matrix:

$$\mathcal{C} = \mathcal{M}_2 \cdot \mathcal{P} \cdot \mathcal{M}_1, \tag{2.3}$$

Here, subscripts 1 and 2 denote the matrices (and corresponding coefficients) associated with the first and second mirrors, respectively, allowing for the possibility of non-identical mirror properties.

Our primary interest lies in the optical response of the cavity, specifically its reflection and transmission characteristics. Effectively, the cavity resonator acts as a mirror-like optical element that can again be treated in the form introduced in Equation (2.1). We therefore define the amplitude reflection and transmission coefficients of the cavity, $r_\mathcal{C}$ and $t_\mathcal{C}$, as:

$$r_\mathcal{C} = -i\, t_\mathcal{C} \cdot \mathcal{C}_{21} = -i\, \frac{\mathcal{C}_{21}}{\mathcal{C}_{22}}, \tag{2.4a}$$

$$t_\mathcal{C} = \frac{1}{\mathcal{C}_{22}}, \tag{2.4b}$$

where $\mathcal{C}_{ij}$ are the entries of the composite cavity matrix $\mathcal{C}$. From these expressions, the intensity reflectance $\mathcal{R}_\mathcal{C}$ and transmittance $\mathcal{T}_\mathcal{C}$ follow directly. Assuming identical mirrors, i.e., $r \equiv r_1 = r_2$ and $t \equiv t_1 = t_2$, we obtain:

$$\mathcal{R}_\mathcal{C} = |r_\mathcal{C}|^2 = \frac{\mathcal{R}}{(1-\mathcal{R})^2} \cdot \frac{(1-\mathcal{R}-\mathcal{T})^2 + 4(\mathcal{R}+\mathcal{T})\sin^2(kd)}{1 + \frac{4\mathcal{R}}{(1-\mathcal{R})^2}\sin^2(kd)}, \tag{2.5a}$$

$$\mathcal{T}_\mathcal{C} = |t_\mathcal{C}|^2 = \underbrace{\frac{\mathcal{T}^2}{(1-\mathcal{R})^2}}_{\eta} \cdot \frac{1}{1 + \frac{4\mathcal{R}}{(1-\mathcal{R})^2}\sin^2(kd)}, \tag{2.5b}$$

where $\mathcal{R} = |r|^2$ and $\mathcal{T} = |t|^2$ denote the intensity reflectance and transmittance of the mirrors. Moreover, the prefactor η sets the contrast of the cavity transmission response. In the absence of non-radiative losses $\mathcal{L}$, i.e., $\mathcal{R} + \mathcal{T} = 1$, the contrast is maximal and

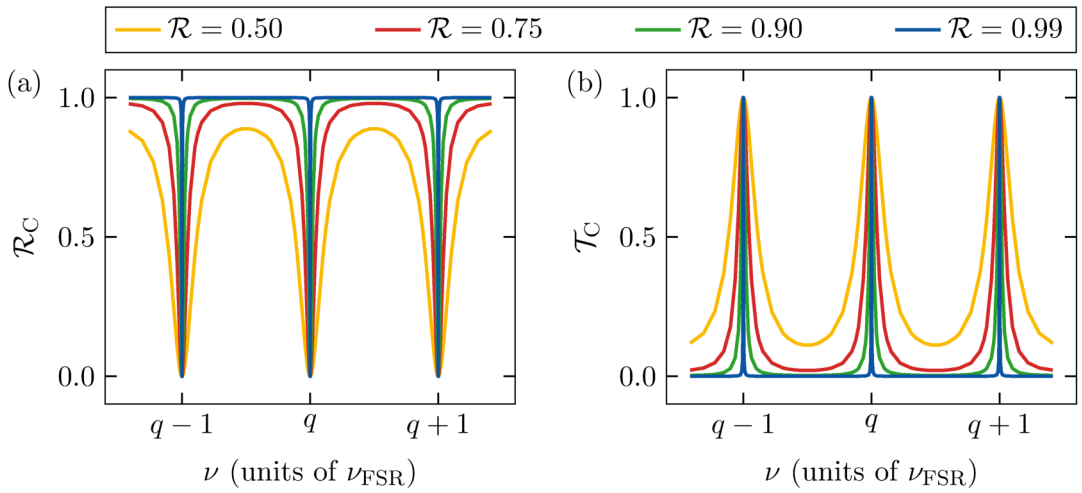


Figure 2.2.: (a) Reflection and (b) transmission spectra of a Fabry–Pérot cavity versus probe frequency ν for different mirror reflectivities $\mathcal{R}$. On resonance, the cavity is highly transmissive, while increasing $\mathcal{R}$ narrows the resonance linewidths, yielding a comb of Lorentzian modes.

decreases with increasing loss[2].

The reflection and transmission spectra of an ideal loss-free optical cavity are shown in Figure 2.2 for different mirror reflectivities $\mathcal{R}$, plotted as a function of the incident light frequency ν. The resulting spectrum exhibits a series of equidistant resonant modes that allow transmission, despite the cavity being composed of highly reflective mirrors.

These resonances occur at wavelengths that satisfy the cavity resonance condition, which requires that an integer multiple of the optical wavelength matches the total round-trip distance within the cavity, $2d$, ensuring constructive interference. Consequently, the frequency spacing between adjacent modes, labeled q and $q \pm 1$, is given by:

$$\begin{aligned} \Delta\nu_q &= |\nu_{q\pm1} - \nu_q| \\ &= \frac{c}{2nd}, \end{aligned} \tag{2.6}$$

and is known as the free spectral range (FSR), ν_{FSR} [43]. This parameter is a fundamental characteristic of the cavity resonator, establishing a direct relationship between the cavity length d and its frequency scale.

[2]Asymmetric mirror transmittances (or, equivalently, unequal input and output coupling rates) also reduce the interference contrast, since the intracavity field cannot build up and outcouple symmetrically. Additional intracavity loss can shift the system towards an optical impedance-matching (critical-coupling) condition, analogous to impedance matching in electrical circuits [42].

2.1.1. Limit of high mirror reflectivities

The presented spectra in Figure 2.2 further demonstrate that as the mirror reflectance approaches unity, the resonant modes become increasingly narrow, resulting in well-separated cavity modes. Under these conditions, the transmission function (Equation (2.5b)) of a single mode with resonance frequency $\nu_q = q \cdot \nu_{\mathrm{FSR}}$ can be approximated by:

$$\mathcal{T}_C(\nu_q) \underset{\nu_q \approx q \cdot \nu_{\mathrm{FSR}}}{\approx} \eta \frac{\left(\nu_{\mathrm{FSR}} \cdot \frac{1-\mathcal{R}}{\pi\sqrt{\mathcal{R}}}/2\right)^2}{\left(\nu_{\mathrm{FSR}} \cdot \frac{1-\mathcal{R}}{\pi\sqrt{\mathcal{R}}}/2\right)^2 + \left(\nu_q - q \cdot \nu_{\mathrm{FSR}}\right)^2}. \tag{2.7}$$

This function corresponds to a Lorentzian lineshape centred at $\nu_q = q \cdot \nu_{\mathrm{FSR}}$ with a full width at half maximum (FWHM) given by:

$$\delta\nu = \frac{\nu_{\mathrm{FSR}}}{\mathcal{F}}, \quad \text{with } \mathcal{F} = \frac{\pi\sqrt{\mathcal{R}}}{1-\mathcal{R}}, \tag{2.8}$$

where $\mathcal{F}$ is the finesse parameter, which relates the width of a single mode to the separation of adjacent modes as a function of the end mirror reflectivity[3]. The finesse serves as second key parameter characterising the optical response of a cavity resonator, quantifying the "sharpness" of the resonances. When combined with the free spectral range, these parameters enable a comprehensive reconstruction of the cavity's reflection and transmission spectra.

In general, the reflectance $\mathcal{R}$, transmittance $\mathcal{T}$, and non-radiative losses $\mathcal{L}$ — such as absorption and scattering — must sum to unity, i.e., $\mathcal{R}+\mathcal{T}+\mathcal{L} = 1$. Using this relation, in the limit of high reflectivities the finesse can be further approximated by [9]:

$$\mathcal{F} \underset{\mathcal{R} \lesssim 1}{\approx} \frac{2\pi}{\sum_i \mathcal{T}_i + \sum_j \mathcal{L}_j}, \tag{2.9}$$

where $\mathcal{T}_i$ and $\mathcal{L}_j$ denote the transmissive and non-radiative loss components associated with each optical element, respectively. Thus, with known mirror transmittances, this relationship allows us to infer intracavity losses. Consequently, the finesse is a particularly critical parameter in the study of dissipative cavity systems.

[3] This definition of the finesse is valid even for low reflectivities. However, in this regime, the comparison of the linewidth to the free spectral range of the spectrum does not strictly equate to the finesse, so that small finesse parameter only qualitatively represents the "sharpness" of the resonance.

2.2. Transversal electromagnetic (TEM) cavity modes

The cavity model discussed thus far assumes two planar mirrors positioned opposite each other, with the circulating field treated as a one-dimensional wave propagating strictly along the optical axis. In practice, however, optical fields possess a finite transverse extent and are more accurately described as beam-like solutions of the wave equation. In the paraxial regime, where the field divergence remains small, the wave equation reduces to the paraxial Helmholtz equation, which provides the formalism for describing the transverse mode structure of optical resonators [43]. These transverse modes impose boundary conditions on the mirror geometry and lead to a stability criterion for resonant field buildup. In addition to the fundamental Gaussian mode, the cavity geometry can also support higher-order solutions with more complex transverse mode profiles, such as Laguerre–Gaussian modes. These modes become relevant for controlling cavity–cavity coupling in Chapter 5.

2.2.1. Fundamental Gaussian beam

The Gaussian beam is a fundamental solution to the paraxial Helmholtz equation, providing a comprehensive framework for describing optical beam propagation in the paraxial regime. In the following, we summarise the key results for Gaussian beam propagation, as presented in Ref. [43]. The complex electric field distribution of a Gaussian beam propagating along the z-axis is given by:

$$\boldsymbol{\mathcal{E}}(z,r) = \boldsymbol{\mathcal{E}}_1 \frac{w_0}{w(z)} \exp\left(-\frac{r^2}{w^2(z)}\right) \exp\left(-ikz - ik\frac{r^2}{2R(z)} + i\psi(z)\right), \tag{2.10}$$

where $\boldsymbol{\mathcal{E}}_1$ represents the polarisation and maximal amplitude of the electric field. Along the longitudinal direction, the beam exhibits characteristics similar to those of a plane wave, with a phase term expressed as $kz + \psi(z)$, where k denotes the wavevector of the optical field. The Gouy phase shift $\psi(z)$ accounts for an additional phase delay in the wavefront compared to a plane wave, particularly in regions far from the beam's focal point [43].

However, due to the finite transverse extent of the beam, interference effects arising from different transverse positions within the xy-plane lead to the formation of curved wavefronts when deviating from the optical axis. These wavefronts are characterised by their reciprocal radius of curvature, $R^{-1}(z)$.

In the transverse direction, the electric field follows a Gaussian profile, exhibiting a bell-shaped distribution centred at $r = 0$, as illustrated in Figure 2.3.(a). The beam width is conventionally defined by the radial position at which the field amplitude decreases to $1/e$ of its maximum value, denoted by $w(z)$. The minimum beam waist w_0 occurs at the focal plane (i.e., $w_0 \equiv w(z = 0)$) and represents the narrowest transverse extent of the beam.

The key parameters describing the Gaussian beam quantitatively are interrelated through the Rayleigh length z_R[4] and are defined as:

$$w_0 = \sqrt{\frac{\lambda_0 z_R}{\pi n}}, \tag{2.11a}$$

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

$$R(z) = z\left(1 + \left(\frac{z_R}{z}\right)^2\right), \tag{2.11c}$$

$$\psi(z) = \arctan\left(\frac{z}{z_R}\right), \tag{2.11d}$$

where λ_0 is the vacuum wavelength of the optical mode and n is the refractive index of the ambient material.

Figure 2.3.(b) illustrates the evolution of these beam parameters along the longitudinal axis, depicting both the variation in mode size $w(z)$ and the curvature of the wavefront $R(z)$ at different propagation distances z. As the distance from the focal plane increases, the beam width exhibits an approximately linear expansion, while the wavefront curvature asymptotically approaches that of a spherical wave, where $R(z) \approx z$. Conversely, at the Rayleigh length ($z = z_R$), the radius of curvature reaches its minimum value, while the mode size increases by a factor of $\sqrt{2}$ relative to the minimum beam waist w_0. In this sense, z_R therefore characterises the divergence properties of the Gaussian beam. Hence, by adjusting the Rayleigh length, a Gaussian beam can describe configurations ranging from strongly focused to highly collimated beams, providing a versatile framework for modelling optical beam propagation while preserving the wave nature of the optical mode.

Notably, Figure 2.3.(a) illustrates the transverse intensity distribution at two distinct positions along the cavity axis, demonstrating that while the transverse profile evolves with z, the qualitative behaviour of the beam can be effectively decomposed into longitudinal and transverse components [42].

2.2.2. Boundary condition of a resonator with curved end mirrors

In an optical cavity resonator, the formation of a resonant mode is governed by interference effects, which can only be sustained if constructive interference is maintained throughout the entire mode volume. Due to the curved phase fronts (wavefronts) of beam-like transverse modes, this requirement is met when the curvature of the reflecting surfaces precisely matches that of the propagating wavefronts. In this scenario, phase

[4] In general, the Gaussian beam is characterised by the q-parameter [43] or complex beam radius [42], defined by $\frac{1}{q(z)} = \frac{1}{R(z)} - 2i\frac{1}{kw^2(z)}$. However, this parameter is connected to the Rayleigh length via $z_R = \mathfrak{Im}(q(z))$. While the q-paramter provides a powerful tool to describe the propagation of the Gaussian beam through an optical system, this analysis is beyond the scope of the thesis. A complete description can be found in literature (e.g., Ref. [43]).

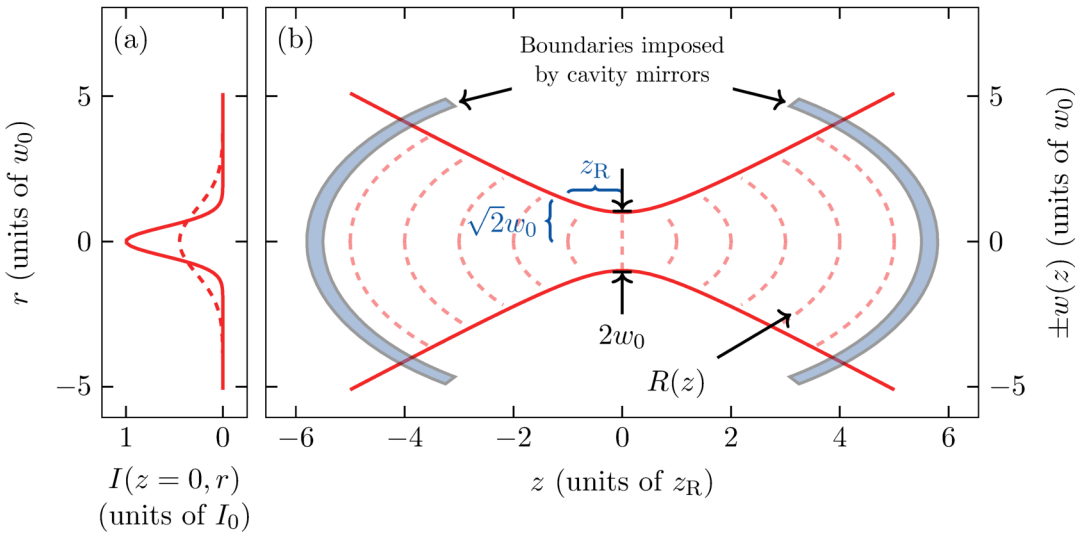


Figure 2.3.: (a) Normalised transverse intensity profile in the focal plane (solid) and at $z = 2z_R$ (dashed), illustrating the Gaussian distribution. Along the optical axis, the intensity profile changes only in width and amplitude. (b) Illustration of the Gaussian beam radius $w(z)$ along the propagation axis, normalised to the beam waist w_0 and Rayleigh length z_R. The waist w_0 and its relation to z_R are indicated. The wavefront curvature $R^{-1}(z)$ is shown at different propagation distances. For stable cavity operation, the mirror curvature must match the wavefront curvature; one possible configuration is indicated by the mirrors sketched in blue.

discrepancies arising from optical path differences between the wavefront and the mirror surface are eliminated, ensuring constructive interference. As a result, the cavity mode, described by the Gaussian beam-like solution of the paraxial Helmholtz equation, is subject to boundary conditions dictated by the radii of curvature of the end mirrors, given by:

$$-R_1 = z_1 + \frac{z_R^2}{z_1}, \tag{2.12a}$$

$$R_2 = z_2 + \frac{z_R^2}{z_2}, \tag{2.12b}$$

$$d = z_2 - z_1, \tag{2.12c}$$

where $z_{1/2}$ denote the positions of the mirrors relative to the beam waist, d represents the total mirror separation, z_R is the Rayleigh length of the cavity mode, and $R_{1/2}$ are the radii of curvature of the left and right mirrors, respectively. A positive values of R_i corresponds to a concave mirror, while a negative value (i.e., $R_i < 0$) describes a convex curvature. Conversely, the radius of curvature of a Gaussian beam changes sign on either side of its focal point. Hence, a relative minus sign was introduced in Equation (2.12a) to correctly describe the boundary conditions [44].

Following Ref. [44], we can use Equation (2.11c) to rearrange the expressions for the mirror positions z_1 and z_2, as well as the Rayleigh length z_R (and consequently the beam waist w_0) and rewrite them in terms of the mirror curvatures and their separation:

$$z_1 = -\frac{g_2(1-g_1)}{g_1+g_2-2g_1g_2}d, \tag{2.13a}$$

$$z_2 = \frac{g_1(1-g_2)}{g_1+g_2-2g_1g_2}d = z_1 + d, \tag{2.13b}$$

$$z_\mathrm{R} = \sqrt{\frac{g_1g_2(1-g_1g_2)}{(g_1+g_2-2g_1g_2)^2}}d, \tag{2.13c}$$

$$w_0 = \sqrt{\frac{\lambda_0 d}{\pi n}}\left(\frac{g_1g_2(1-g_1g_2)}{(g_1+g_2-2g_1g_2)^2}\right)^{\frac{1}{4}}. \tag{2.13d}$$

Here, the mirror curvatures and cavity length are conveniently summarised in two dimensionless quantities, the g-parameters of the mirrors, defined as:

$$g_i = 1 - \frac{d}{R_i}. \tag{2.14}$$

These equations highlight that the optical mode with wavelength λ_0/n is fully determined by the g_1- and g_2-parameters of the mirrors, as well as their separation d. Notably, the imposed boundary conditions allow for at most one valid solution for any given cavity geometry. This implies that if a stable optical mode exists, it is uniquely associated with a specific Rayleigh length.

Stability criterion

Examining Equation (2.13c), it becomes evident that not all combinations of g_1 and g_2 yield a real-valued Rayleigh length. This inconsistency with the Rayleigh length's definition signifies an unstable cavity mode, in which the boundary conditions cannot be satisfied [43]. Physically, in such cases, successive reflections of the circulating light beam fail to establish a closed optical path, preventing the formation of a standing wave. As a result, the mode undergoes destructive interference or experiences transverse drift, ultimately escaping the cavity volume [42, 44].

From Equation (2.13c), the requirement for a real-valued Rayleigh length z_R is given by the inequality:

$$g_1g_2(1-g_1g_2) > 0. \tag{2.15}$$

Rearranging this expression leads to the confinement criterion for a stable optical cavity resonator [43]:

$$0 \leq g_1g_2 \leq 1. \tag{2.16}$$

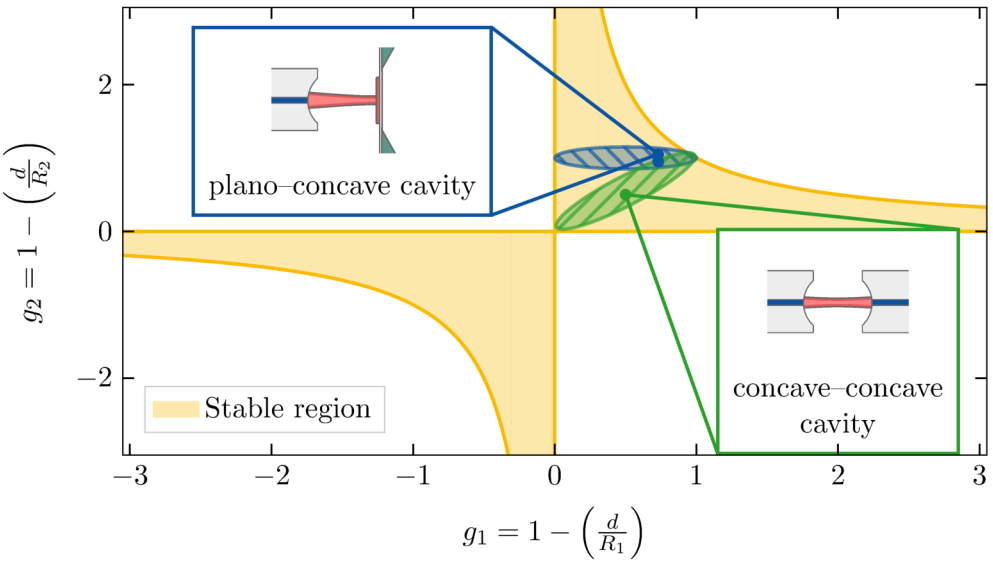


Figure 2.4.: Graphical representation of the stability regions of a two-mirror cavity with mirror separation d, parameterised by the dimensionless mirror parameters g_1 and g_2. For parameter pairs (g_1, g_2) within the orange shaded region, the stability criterion is satisfied and a stable cavity mode exists. The green and blue shaded regions indicate typical parameter ranges for plano–concave and symmetrically concave cavities, respectively, as used in this thesis.

This stability condition is visualised in Figure 2.4, where g_1 and g_2 are plotted against each other. The orange-shaded region represents the parameter space where the inequality holds true, corresponding to stable cavity mode configurations.

In this thesis, we primarily investigate two coupled plano-concave cavities (with $g_2 \approx 1$), formed between a planar coupling mirror and a concave end mirror. However, the cavities are formed by inserting a plane mirror into a pre-constructed symmetrically concave cavity resonator (with $g_1 \approx g_2$). The pre-constructed cavity is subject of Section 2.3, and the plano-concave cavities are established in Chapter 3 following further investigation in subsequent chapters. As reference, both cavity geometries are also sketched in Figure 2.4, and the approximate regimes are highlighted in the parametric representation of the stability criterion in green and blue colours, respectively.

2.2.3. Higher order transversal electromagnetic wave solutions of an optical cavity resonator

The Gaussian beam is not the only solution to the paraxial Helmholtz equation; rather, it represents just one mode within a broader set of eigenmodes supported by an optical resonator. In fact, the paraxial Helmholtz equation admits a manifold of solutions, each corresponding to a distinct transverse mode of the cavity. To establish a mathematical and conceptual foundation for the analysis in Chapter 5, we consider one particular class of cavity eigenmodes that extends beyond the fundamental Gaussian mode.

To describe these higher-order modes, it is convenient to express the cavity eigenmodes in cylindrical coordinates, leveraging the axial symmetry of the system. In this representation, the general solutions to the paraxial Helmholtz equation take the form:

$$\begin{aligned}\boldsymbol{\mathcal{E}}_{lp}(z,r) =& \boldsymbol{\mathcal{E}}_{1,lp}\frac{w_0}{w(z)}\left(\frac{r}{w(z)}\right)^l L_p^l\left(\frac{2r^2}{w^2(z)}\right)\exp\left(-\frac{r^2}{w^2(z)}\right) \\ &\times\exp\left(-ikz - ik\frac{r^2}{2R(z)} - il\phi + i(l+2p+1)\psi(z)\right),\end{aligned} \tag{2.17}$$

where $\boldsymbol{\mathcal{E}}_{1,lp}$ again describes the polarisation and maximal amplitude of the electric field distribution for a given mode, while w_0, $w(z)$, $R(z)$, and $\psi(z)$ are still defined by Equations (2.11a) to (2.11d), such that the longitudinal evolution qualitatively follows that of the Gaussian beam. The indices l and p characterise the order of the mode, and the functions L_p^l correspond to the generalised Laguerre polynomials, modifying the transverse mode structure [43, 45].

These polynomials also appear in the quantum mechanical description of the hydrogen atom, suggesting a mathematical analogy between the mode indices l and p in this context and the principal quantum number n and angular momentum quantum number l of an electron bound to a proton [46]. This analogy becomes particularly evident when examining the transverse intensity distributions of these modes. Figure 2.5 illustrates the normalised intensity profiles — proportional to the absolute square of the electric field distribution — for the first nine modes, corresponding to $l, p \in \{0, 1, 2\}$.

The integer p determines the number of radial intensity minima, analogous to the nodes in the electric field, much like the quantum number $(n - l - 1)$ corresponds to the number of nodes in the wave function of the electron. Meanwhile, increasing l results in a broader spatial extent of the mode, shifting the intensity away from the optical axis. For $l = 0$, the mode exhibits maximum intensity at the beam centre, whereas for $l > 0$, the central intensity vanishes, forming a ring-shaped intensity profile. This behaviour closely resembles the distinction between states with and without angular momentum in the atomic counterpart, where nonzero angular momentum states exhibit a node at the centre, demonstrating the absence of an electron probability density at the nucleus [46].

The phase evolution of the electric field along the optical axis is governed by the azimuthal phase term in Equation (2.17), which is proportional to $l\phi$. This phase dependence results in helical wavefronts that rotate around the optical axis at a rate proportional to l. Physically, this structure directly implies an associated orbital angular momentum of the optical mode [43].

It is important to note that Laguerre–Gaussian modes form a complete orthonormal basis for all solutions of the paraxial Helmholtz equation in a cylindrically symmetric optical cavity. However, an alternative decomposition can be obtained by expressing the field in cartesian coordinates (x, y, z), which yields Hermite–Gaussian modes [43]. These modes also constitute a complete and orthonormal basis set. Notably, by superimposing modes from one representation, it is possible to reconstruct modes from the other [42], demonstrating the fundamental equivalence of these mathematical descriptions.

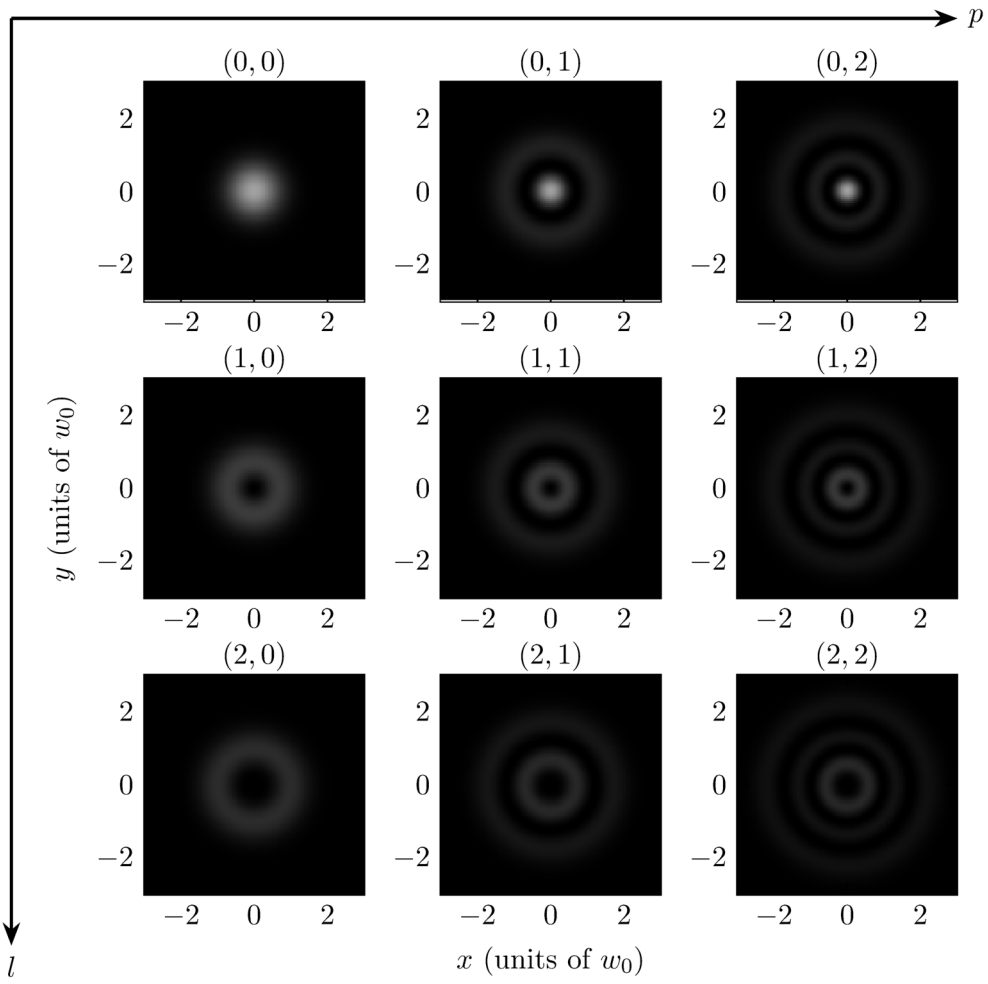


Figure 2.5.: Normalised intensity profiles of higher-order Laguerre–Gaussian modes supported by a cylindrically symmetric cavity, shown in the plane of the beam waist. The intensity is normalised to the peak intensity of the fundamental mode, with $(l, p) = (0, 0)$. The azimuthal and radial mode indices l and p increase from top to bottom and from left to right, respectively.

2.3. Experimental realisation of a fibre-based cavity apparatus

The coupled cavity setup realised in this thesis is based on a microcavity formed between the end facets of two machined single-mode optical fibres [9]. The fabrication process of these fibre mirrors involves laser ablation to imprint the desired laser profile onto the fibre end facets, creating a concave indent, followed by the application of a high-reflectivity mirror coating. The intrinsic fibre coupling of the fabricated mirrors then enables highly efficient detection of dissipative processes, as the leakage field escaping from the cavity can be directly measured.

Furthermore, the combination of large mirror curvature, compact cavity lengths, and

strongly confined mode sizes alleviates the constraints on the membrane mirror — analysed in Chapter 3 — by permitting larger mirror curvatures and smaller mirror sizes. This aspect is particularly significant in the context of the cavity stability criterion (see Section 2.2.2), where the mirror geometry plays a crucial role in maintaining stable cavity operation. These advantages establish the fibre-based microcavity as an ideal platform for initial investigations of a coupled cavity system, providing a robust and accessible testbed for exploring fundamental mechanisms in coupled dissipative quantum systems.

The fibre mirrors used in this thesis were fabricated by Pia Fürtjes as part of her master's thesis in 2018, within the framework of a collaborative project between the *Institut für Angewandte Physik* (IAP) and the *Physikalisches Institut* (PI) at the University of Bonn [10]. This project focused on developing in-house fibre mirrors, primarily for applications in quantum technologies, including quantum repeaters and miniaturised spectroscopy systems [47–50]. A detailed description of the fabrication process is available in her thesis [51] as well as in subsequent studies [52, 53].

To ensure optimal functionality of the fibre cavity within the coupled system, we perform a preliminary analysis of the cavity resonator formed between the fibre mirrors. This allows us to assess the alignment quality of the fibre cavity by examining transverse mode effects as a function of cavity length [40]. This intermediate step is particularly useful, as geometric imperfections — such as deviations in centring or tilt in the imprinted mirror profile — manifest as characteristic signatures in the coupling efficiency and finesse of the cavity resonator. The extracted reference parameters provide a reliable foundation for subsequent experimental decisions, ultimately guiding the development of the final coupled cavity configuration.

2.3.1. Measurement of the reflection spectrum of the fibre cavity

To characterise the fibre mirrors, we first assemble a fibre-based Fabry–Pérot cavity in a dedicated test setup. One fibre is securely fixed with its coated end facet suspended in air, while the opposing fibre is mounted on a five-axis translation stage. This stage allows precise positioning along the x-, y-, and z-axes as well as angular alignment of the fibre tips. Piezoelectric actuators integrated into the stage provide fine modulation of the mirror separation and thus enable controlled scans of the cavity resonance when driven by an arbitrary waveform generator (AWG).

Initial alignment and a coarse estimate of the cavity length are obtained using two USB microscopes mounted perpendicular to each other, providing visual feedback in both transverse directions. A representative microscope image, together with a schematic of the optical setup, is shown in Figure 2.6.

For the optical characterisation, we probe the cavity using a diode laser at a wavelength of $\lambda \approx 780$ nm. The incident field is phase-modulated by an electro-optic modulator (EOM) driven at radio frequency Ω, generating frequency sidebands at detuning $\pm\Omega$ with respect to the carrier. These sidebands provide well-defined frequency markers that are later used to calibrate cavity-length scans in units of optical detuning.

After passing a beam splitter, the probe field is coupled into the input fibre. The light reflected from the cavity is detected on photodiode $\mathrm{PD}_{\mathcal{R}}$, while the light transmitted

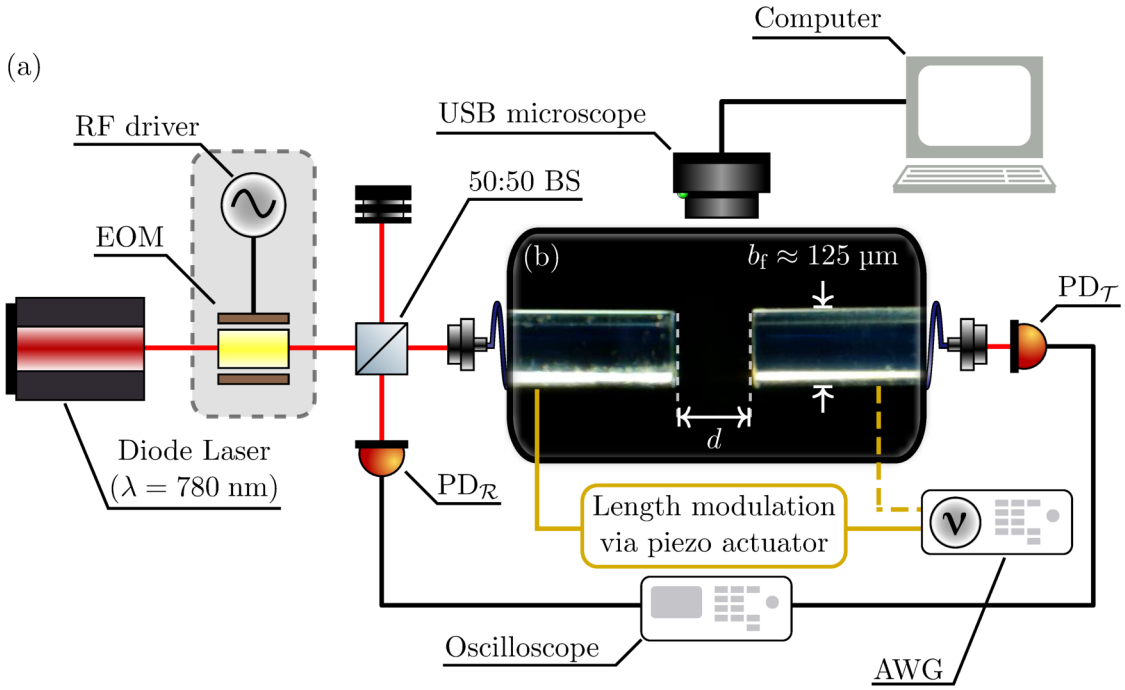


Figure 2.6.: (a) Optical setup used to characterise the fibre-based cavity resonator. The cavity is probed with a diode laser, and the reflected and transmitted signals are detected on photodiodes. An electro-optic modulator (EOM) can be used to generate sidebands that serve as frequency markers. The cavity length is scanned by driving a piezo actuator with an arbitrary waveform generator (AWG). (b) USB microscope image used to estimate the cavity length using the fibre diameter as a length reference. This figure includes elements generated with Ref. [54].

through the opposing fibre is monitored on photodiode $PD_{\mathcal{T}}$.

The measured reflection spectrum, acquired with and without EOM modulation, is shown in Figure 2.7.(a) and Figure 2.7.(b), respectively. The signal recorded on $PD_{\mathcal{R}}$ is normalised to an off-resonant reference measurement. By fitting the resonance with a Lorentzian lineshape (red), we extract the coupling strength to the cavity mode, η, and determine the cavity linewidth $\delta\nu$. The finesse is then obtained from:

$$\mathcal{F} = \frac{\nu_{\mathrm{FSR}}}{\delta\nu}. \tag{2.18}$$

In the raw data, the horizontal axis corresponds to the time elapsed after the AWG trigger and therefore does not directly provide a frequency scale. We calibrate this axis using the EOM sidebands, whose detuning $\pm\Omega$ imposes a known frequency separation on the reflected spectrum, as illustrated in Figure 2.7.(b).

To evaluate ν_{FSR}, an approach based on scanning an entire spectral period is impractical, as the piezo response becomes increasingly nonlinear at larger displacements and would compromise calibration accuracy. Instead, we estimate the cavity length directly

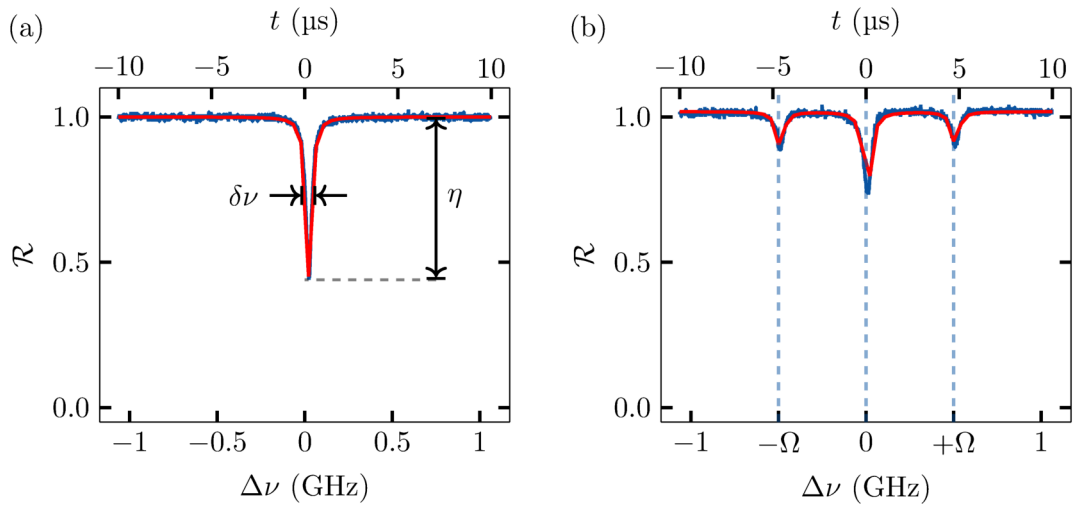


Figure 2.7.: Normalised reflection spectrum of a fibre–fibre cavity as a function of the detuning $\Delta\nu$ between the carrier and the cavity resonance frequency. The probe field consists of (a) the carrier only or (b) the carrier accompanied by two sidebands detuned by $\pm\Omega$. The sidebands provide frequency markers that calibrate the cavity-length scan. Subsequently, the coupling efficiency η and the linewidth $\delta\nu$ are extracted from the carrier resonance.

from microscope images by referencing the well-defined fibre diameter $b_\mathrm{f} = 125(1)$ µm [55]. A representative image used for this estimate is shown in Figure 2.6.(b), where both the fibre width and the mirror separation are annotated. From the resulting cavity length d, we compute ν_FSR and thereby obtain the finesse $\mathcal{F}$.

For the measurement shown in Figure 2.7, we determine a cavity length of $d = 111(9)$ µm, corresponding to a finesse of $\mathcal{F} \approx 26(2) \times 10^3$. This confirms operation in the high-finesse regime, characterised by strong field confinement and narrow resonance features.

2.3.2. Limitations due to transverse cavity mode effects

The reflection spectrum of the cavity at a fixed cavity length provides only a static representation of the cavity alignment. To obtain a more complete characterisation, we repeat the described measurement procedure over a range of fibre separations and extract the coupling strength η and finesse $\mathcal{F}$ as functions of the cavity length d. As the cavity length increases, the transverse cavity mode undergoes modifications, introducing characteristic signatures in both η and $\mathcal{F}$. By analysing these effects, we identify the intrinsic limitations of fibre-based microcavities.

Furthermore, by comparing our experimental results to established theoretical models describing transverse mode effects [9, 56], we assess how closely the system operates to its optimal configuration. This comparison provides an indirect yet meaningful measure of cavity performance. To ensure an accurate representation of the transverse alignment of the optical axis with respect to the fibre tips, we avoid angular realignments while

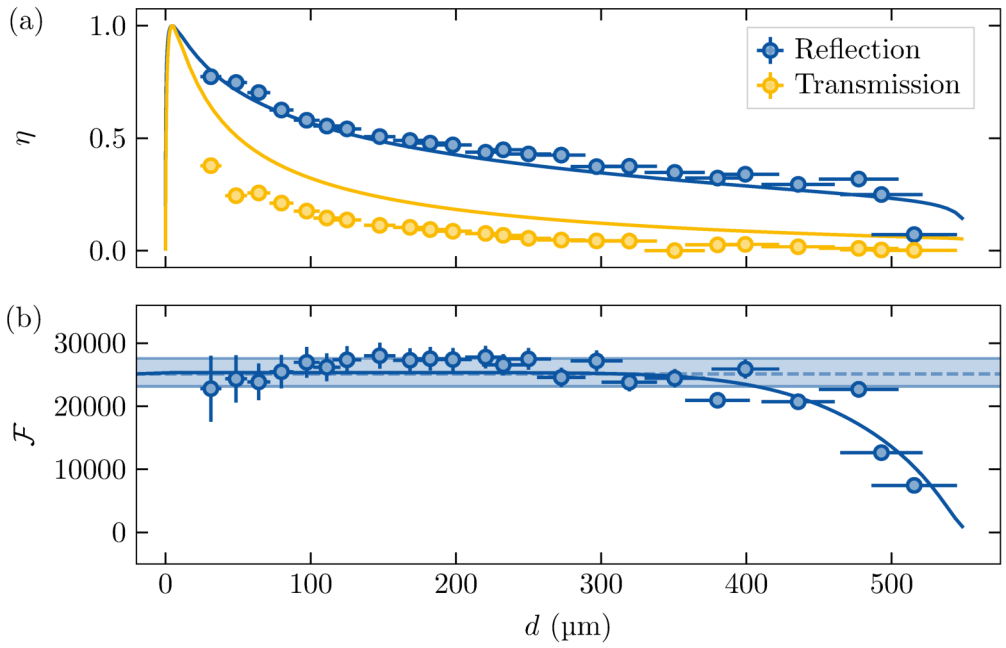


Figure 2.8.: (a) Coupling efficiency η and (b) finesse $\mathcal{F}$ versus fibre separation d. The agreement with the models (solid lines) indicates good alignment of the input coupling; the reduced transmission contrast (orange) suggests a transverse offset at the output. The drop in $\mathcal{F}$ at large d is attributed to clipping losses due to the expanding cavity mode.

extending the cavity length. However, as the z-axis of the translation stage and the cavity axis are generally not perfectly aligned, small compensatory adjustments in the xy-plane are applied to maximise the observed coupling strength. The results of this analysis are summarised in Figure 2.8.

Reduction in coupling strength due to reduced mode matching

The coupling efficiency between the incident field and the cavity mode is determined by their transverse mode overlap. In fibre-based Fabry–Pérot cavities, the incident field is fixed by the guided fibre mode, which limits the attainable overlap with the cavity eigenmode [9]. Since the cavity mode size at the mirrors varies with cavity length d, the coupling efficiency η is inherently d-dependent. Characterising this dependence allows us to identify cavity lengths that maintain sufficient transverse access while preserving efficient coupling. Moreover, systematic deviations from the expected behaviour can serve as a diagnostic for imperfect alignment [57] or additional intracavity loss channels (see Equation (2.5b)).

The measured coupling efficiency as a function of cavity length is shown in Figure 2.8.(a). The coupling is extracted from the optical contrast in reflection (blue) and transmission (orange), with all values normalised to the off-resonant reflection level. The transmission signal is additionally corrected for differences in detection efficiencies,

as the reflected signal passes the beam splitter one additional time.

The coupling efficiency can be modelled by the overlap integral between the fibre mode $\boldsymbol{\mathcal{E}}_{\mathrm{f}}(r)$ and the transverse cavity mode $\boldsymbol{\mathcal{E}}_{1/2}(z_{1/2}(g_1, g_2, d), r)$ evaluated at the mirror positions $z_{1/2}(g_1, g_2, d)$:

$$\begin{aligned}\eta_{1/2}(d) &= \frac{\int \mathrm{d}S\ \boldsymbol{\mathcal{E}}_{1/2}(z_{1/2}(g_1, g_2, d), r)\boldsymbol{\mathcal{E}}^{\star}_{\mathrm{f}}(r)}{\left(\int \mathrm{d}S\ |\boldsymbol{\mathcal{E}}_{\mathrm{f}}(r)|^2 \times \int \mathrm{d}S\ |\boldsymbol{\mathcal{E}}_{1/2}(z_{1/2}(g_1, g_2, d), r)|^2\right)^{\frac{1}{2}}} \\ &\approx \frac{4}{\left(\frac{w_{\mathrm{f}}}{w_{1/2}(g_1,g_2,d)} + \frac{w_{1/2}(g_1,g_2,d)}{w_{\mathrm{f}}}\right)^2},\end{aligned} \tag{2.19}$$

where $w_{1/2}$ denotes the cavity mode radius at the fibre mirror positions $z_{1/2}$ (see Equations (2.13a) to (2.13d)). The transverse cavity mode profile is determined by the imprinted mirror curvatures, which are set during fibre fabrication. The parameters used for the model are obtained from independent fibre-mirror characterisation [51] and are summarised in Table 2.1.

The guided fibre mode is approximated as Gaussian with a mode diameter of $2w_{\mathrm{f}} \approx 6$ µm [55]. As it propagates in the waveguide, the fibre mode is taken to exhibit negligible wavefront curvature over the relevant interaction region. The resulting theoretical prediction for $\eta_1(d)$ is shown as a solid line in Figure 2.8.(a). For transmission, the expected coupling scales as $\eta_1(d) \cdot \eta_2(d)$, accounting for mode matching at both the input and output fibre interfaces [9].

The experimental results show excellent agreement with the theoretical predictions in the reflection data, suggesting a well-aligned mode between the cavity and the input fibre. Within the sensitivity of this measurement, the observed reduction in coupling is consistent with mode mismatch alone, suggesting that additional loss mechanisms play no dominant role (see Equation (2.5b)). However, the optical contrast in the transmission signal is consistently overestimated by the model, indicating a potential systematic misalignment between the cavity mode and the output fibre mode. The constant factor of two-fold reduction in transmission coupling across all cavity lengths suggests that the misalignment is length-independent, likely arising from a transverse offset between the cavity and fibre mode centres.

Increasing losses due to mode clipping

Figure 2.8.(b) shows the measured finesse $\mathcal{F}$ as a function of cavity length d. In Section 2.1.1, the finesse was introduced as a length-independent parameter, and therefore it is expected to remain approximately constant as long as the mirror losses do not change. Over most of the investigated range, the data are consistent with this behaviour, aside from a small local maximum around $d \approx 150$ µm, where the cavity alignment was initially optimised. For cavity lengths exceeding $d \approx 400$ µm, however, the finesse decreases noticeably, indicating increased intracavity losses.

A plausible origin of this degradation is mode clipping at the fibre mirrors. As the

cavity length increases, the transverse cavity mode expands, and beyond a critical length a non-negligible fraction of the optical power extends beyond the effective mirror surface. Importantly, the relevant scale is not necessarily given by the physical fibre end facet: deviations from an ideal parabolic mirror profile can introduce additional loss already at smaller mode diameters. Such deviations arise because the imprinting laser has a Gaussian intensity profile, producing an approximately parabolic curvature only near the mirror centre [9, 51–53]. Additional deviations from the ideal shape can occur due to nonlinear absorption during fabrication, where regions exposed to higher intensities undergo disproportionate material removal [52].

The clipping loss $\mathcal{L}_{\text{clip}}$ at a given mirror can be expressed as the fraction of the transverse mode intensity that lies outside a mirror of effective radius a. For a Gaussian mode, this loss can be approximated as [56]:

$$\mathcal{L}_{\text{clip}} = 1 - \frac{\int_0^a \mathrm{d}r\, r|\boldsymbol{\mathcal{E}}_{1/2}(z_{1/2}, r)|^2}{\int_0^\infty \mathrm{d}r\, r|\boldsymbol{\mathcal{E}}_{1/2}(z_{1/2}, r)|^2} \approx \exp\left(-\frac{a^2}{w_{1/2}^2}\right), \tag{2.20}$$

where $w_{1/2}$ is the mode radius at the corresponding mirror position $z_{1/2}$, as defined in Equations (2.13a) to (2.13d). To capture the length dependence of the finesse, we fit Equation (2.9), i.e.:

$$\mathcal{F} = \frac{2\pi}{2\mathcal{T} + \mathcal{L}_{\text{clip}}}, \tag{2.21}$$

to the data, assuming equal mirror transmittances $\mathcal{T}$ for both fibres and including clipping losses at a single mirror. The latter assumption is motivated by the loss asymmetry discussed in Section 3.3, which indicates that losses at the first (left) mirror dominate. The good agreement between the measured input coupling efficiency and the mode-overlap model (see Figure 2.8.(a)) suggests that additional loss channels do not significantly affect the present measurement.

The resulting fit, shown as a solid line in Figure 2.8.(b), yields an effective mirror radius of $a_1 = 34(1)$ µm. The same model gives a mirror transmittance of $\mathcal{T} = 124(11)$ ppm, corresponding to a high-finesse value of $\mathcal{F} = 25(2) \times 10^3$ (shaded region in the figure). Overall, the agreement with the experimental data supports the interpretation that clipping-induced losses dominate the reduction in finesse at extended cavity lengths and therefore impose a practical upper bound on the accessible cavity length for a given mirror size.

2.3.3. Summary

Taken together, these measurements confirm the successful implementation of a fibre-based microcavity using in-house fabricated fibre mirrors [51]. The key parameters characterising the cavity are summarised in Table 2.1 and serve as a reference for subsequent cavity analyses.

Over the investigated operating range, the cavity exhibits negligible loss, as indicated

Table 2.1.: Summary of the cavity parameters. Length-dependent quantities are evaluated at $d \approx 300$ µm, corresponding to the subsequently fixed cavity length. This characterisation was performed prior to fixing the fibres; changes in the final coupled-cavity implementation are discussed throughout Section 4.2.

Parameter	Value	Unit
R_1	555(6)	µm
R_2	561(17)	µm
d	300(15)	µm
$g_1 g_2$	0.21(1)	–
ν_{FSR}	500(25)	GHz
$\delta\nu$	20(2)	MHz
$\mathcal{F}$	25(2)	10^3
$\mathcal{T}$	124(11)	ppm
$\eta_{\mathcal{R}}$	0.373(1)	–
$\eta_{\mathcal{T}}$	0.0434(1)	–
a_1	34(1)	µm

by the measured coupling efficiency and finesse as functions of cavity length. The observed reduction in coupling efficiency η is fully consistent with mode mismatch between the cavity and fibre modes, indicating that additional loss mechanisms do not play a dominant role. Under these conditions, the measured finesse of $\mathcal{F} = 25(2) \times 10^3$ is set primarily by the finite mirror transmittance, estimated as $\mathcal{T} = \pi/\mathcal{F} = 124(11)$ ppm.

For cavity lengths $d > 400$ µm, a rapid increase in optical losses is observed, which can be attributed to the expanding transverse mode size of the cavity mode. Beyond a critical length, the mode extends beyond the effective mirror area, leading to clipping losses. However, in the present system, the mirror sizes are sufficiently large such that clipping losses become significant only for $d \approx 450$ µm. Furthermore, at mirror separations exceeding $d > R_1 = 555(6)$ µm, the stability criterion (see Equation (2.16)) dictates that cavity formation is no longer possible.

Overall, the agreement between measured and predicted coupling efficiencies indicates a high level of alignment between the fibre and cavity modes. This level of alignment enables a passively aligned and extendable microcavity geometry utilising glass ferrules, in which the fibres can be retracted without requiring active realignment. The resulting transverse access provides sufficient clearance to safely insert a membrane into the cavity volume. The integration of a membrane will be the focus of Chapter 4.

3. Production and characterisation of a dielectric membrane mirror

Contents

3.1. Geometry of the membrane mirror 26
3.1.1. Interferometry of the membrane surface 27
3.1.2. Modifications to the membrane geometry 28
3.2. Optimisation of the mirror coating composition 33
3.2.1. Refractive index of silicon nitride 33
3.2.2. Optimised coating composition 36
3.3. Optical test cavity with a membrane mirror 37
3.3.1. Absorption within the silicon nitride layer 40
3.4. Investigation of the membrane's vibrational modes 42
3.4.1. Detection of the membrane motion 43
3.4.2. Mechanical model 46
3.5. Summary 50

A key challenge in realising coupled-cavity systems in mirror-based micro-architectures is the development of a partially transmitting mirror that divides a single Fabry–Pérot resonator, here formed by two opposing fibre mirrors, into two high-finesse cavities coupled via transmission through this shared element. This experimental configuration imposes stringent requirements on the shared mirror: since both sides of the mirror are embedded within an optical microcavity, it cannot be supported by a conventional macroscopic substrate. Instead, the reflective film must be suspended on a thin carrier membrane stabilised by a surrounding frame. The primary challenge in fabricating such membrane-based mirrors is therefore mitigating the mechanical stress introduced by the high-reflectivity coating [58], which can distort the membrane surface and degrade its optical performance.

This chapter presents how a commercially available silicon nitride (Si_3N_4) membrane can be refined to create a membrane mirror suitable for high-finesse applications. In this context, we demonstrate how modifying the membrane geometry improves surface flatness while preserving high reflectivity, by testing and comparing different membrane shapes. We further characterise the optical properties of both the silicon nitride membrane and the final prototype membrane mirror by performing single-pass transmission spectroscopy. Despite achieving high reflectivity, measurements in a cavity environment reveal noticeable absorption losses within the silicon nitride layer, leading to nonlinear

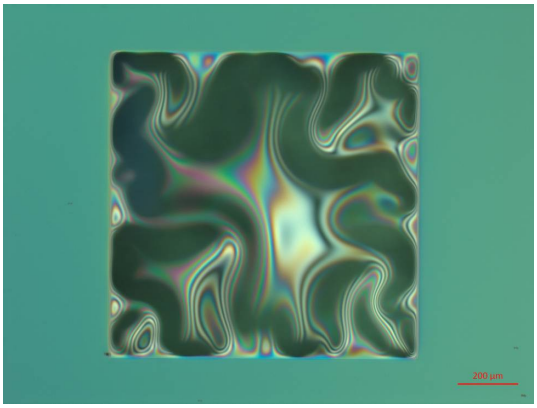


Figure 3.1.: Photography of a coated 1 mm × 1 mm silicon nitride (Si_3N_4) membrane with thickness of 500 nm. The membrane "crumbled" under the internal stress of the applied film, but did not collapse. The image was provided by Laseroptik GmbH.

optical behaviour [59, 60], a phenomenon further discussed in Chapter 6. Finally, we complete the characterisation of the membrane mirror by investigating its vibrational modes in a cavity setup, establishing the basis for possible future optomechanical experiments [30, 61, 62].

3.1. Geometry of the membrane mirror

In our earlier implementation of a "free-standing" mirror, a reflective coating was deposited onto a gold substrate with intentionally poor adhesion, allowing the film to be released after fabrication [63]. While this approach enabled detached mirror films, it introduced several practical limitations: the high internal stress of the coating resulted in pronounced surface curvature, restricting the usable mode size on the mirror and limiting the achievable mode volume. Moreover, the detached films were brittle and difficult to handle. Transfer into a holder typically required a transport liquid (e.g., high-concentration acetone), which frequently contaminated the samples, such that identifying clean mirrors with suitable curvature required repeated sample exchange. This in turn necessitated a removable mount with substantial angular tunability; however, the resulting mechanical flexibility made the setup susceptible to vibrations.

In this chapter, we address these challenges by developing membrane-based mirrors that provide reproducible surface quality while remaining embedded in a mechanically stable mounting structure. The underlying concept is inspired by the "membrane-in-the-middle" configuration [21, 64], a well-established platform in optomechanics for studying interactions between light and mechanical oscillators [30]. In this approach, a thin mechanical membrane oscillator — commonly a commercially available silicon nitride membrane (e.g., from Norcada Inc.[1]) [65] — is positioned near the centre of an optical

[1]Norcada Inc. is a company based in Edmonton, Alberta, Canada that specialises in micro-electromechanical systems. Contact: Eric Daigle (eric@norcada.com)

cavity. Such membranes are typically fabricated by depositing a thin silicon nitride film onto a silicon frame and subsequently removing the underlying silicon by chemical etching, leaving a suspended membrane window [66].

Building on this framework, we transform the membrane from a passive mechanical element into a high-reflectivity mirror suitable for high-finesse cavities by applying a dielectric mirror coating to the thin silicon nitride window [24, 25]. This enables a "mirror-in-the-middle" configuration [20], in which the membrane acts as a functional cavity mirror rather than a weakly reflective dispersive element. However, the high-reflectivity coating introduces substantial stress-driven deformations that degrade the optical performance, as illustrated in Figure 3.1. Annealing is a commonly employed method for relieving internal stress in multilayer films [67, 68]. In our case, however, it did not yield measurable improvements in surface quality. We therefore pursue alternative strategies to mitigate coating-induced curvature that involve adopting symmetric membrane geometries and adding compensatory stress prior the coating deposition.

3.1.1. Interferometry of the membrane surface

The membrane mirror's quality is determined by its radius of curvature and the useable mirror size, i.e., the size of the parabolic profile. In order to assess these quantities and compare the effects of different design modifications, we measure the surface profile using interferometry. For this purpose, we employ an interferometric setup originally developed by Moritz Scharfstädt for the characterisation of fabricated fibre mirrors [53]. The setup uses a green LED as the illumination source, providing a practical compromise between the high axial resolution of white-light interferometry and the long coherence length of monochromatic interferometry. This choice enables fast and robust fringe visualisation across the full membrane surface while maintaining sufficient spatial resolution [69, 70].

The interferometric setup is shown in Figure 3.2.(a). An expanded beam illuminates a Mirau objective [69], which focuses the light onto the membrane surface and simultaneously acts as the interferometric element (detailed in Figure 3.2.(b)). Inside the objective, a 50:50 beam splitter (BS) separates the incident field into a reference beam, reflected from an internal reference surface, and a probe beam, reflected from the target surface. The two beams are recombined at the BS, and their optical path difference produces interference fringes in the detected intensity distribution. The resulting spatially resolved fringe patterns are recorded by a camera. A selection of representative images for different membrane prototypes is shown in Figure 3.3.

While the fringe pattern provides a qualitative indication of the surface profile, it does not unambiguously determine the sign of the curvature (concave versus convex), and contrast variations can be affected by reflectivity inhomogeneities and measurement noise. To obtain a quantitative and unambiguous surface reconstruction, we therefore perform a phase-shifting measurement in which the Mirau objective is displaced along the optical axis using a piezoelectric actuator. This scan changes the optical path length of the probe beam while leaving the reference arm unchanged, thereby modulating the interference phase. Images are acquired at six predefined objective positions, enabling the extraction of the relative phase at each camera pixel [70].

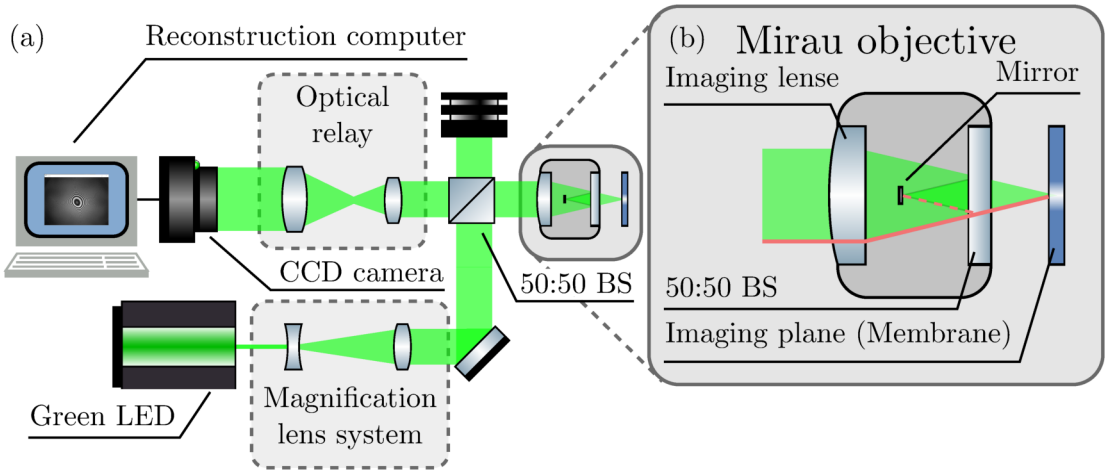


Figure 3.2.: (a) Sketch of the surface interferometry setup, based on a green broadband LED, originally established for reconstruction of shot fibre end facets [53]. We use the setup to asses the surface quality of the (un)coated silicon nitride membranes. (b) Detailed view of the Mirau objective, enabling the interferometry by splitting the incident field in probe and reference beams. Variations in the path length of the probe beam modulate the interference contrast. This figure includes elements generated with Ref. [54].

The resulting phase map $\varphi(x, y) \in [0, 2\pi)$ is subsequently unwrapped to remove 2π discontinuities. Phase unwrapping assumes a sufficiently smooth surface such that the phase difference between neighbouring pixels remains below 2π. Using the centre wavelength of the illumination source ($\lambda = 532$ nm), the relative height profile $\Delta z(x, y)$ is then reconstructed from the unwrapped phase. A detailed description of the optical implementation and the reconstruction algorithm can be found in Refs. [53, 69, 70].

3.1.2. Modifications to the membrane geometry

To mitigate the stress-induced deformation introduced by the high-reflectivity coating, we systematically vary three design parameters and evaluate their impact on the membrane surface profile. First, we investigate the effect of adding a compensatory high-stress silicon nitride layer before applying the coating to the membrane, in order to counteract the coating-induced bending. Second, we study how a circular membrane geometry and reduced lateral dimensions suppress asymmetric deformation patterns and yield a more uniform surface profile. Finally, we combine these approaches and identify an optimised membrane design, which is subsequently characterised in terms of curvature and usable mirror area.

The resulting membrane prototypes are presented in Figure 3.3, showing interferograms that were recorded with the green-light interferometer introduced in the previous subsection. The prototypes are arranged such that neighbouring images differ by only a single design modification, enabling a direct comparison of individual parameters. In

these interferograms, the fringe spacing corresponds to a characteristic height variation on the order of the illumination wavelength, providing an intuitive visualisation of the surface profile.

All membrane mirror variants shown in Figure 3.3 were coated simultaneously by Laseroptik GmbH[2] and therefore accommodate identical mirror coatings. This approach minimises iterative coating cycles and ensures direct comparability between prototypes by eliminating variations in coating design and processing conditions.

Introduction of a high-stress layer

The purpose of the compensatory high-stress layer is to counteract the bending introduced by the high-reflectivity mirror coating by imposing an additional stress contribution of opposite sign. This is achieved by adding a second silicon nitride layer to the membrane, where mechanical interactions at the interface generate an internal stress distribution that partially compensates the deformation. Although both layers consist of silicon nitride, they are deposited under different growth conditions, leading to variations in density and mechanical properties [71] and enabling controlled tuning of the resulting stress.

Based on prior measurements performed by Laseroptik GmbH, the stress introduced by the mirror coating was estimated as $\sigma \approx 450$ MPa. This estimate was obtained from the mechanical response of a coated reference substrate with well-characterised material properties [72]. Using this information, Norcada Inc. designed an appropriate high-stress layer intended to compensate the coating-induced deformation.

The qualitative effect of the compensatory layer is visible in Figure 3.3 when comparing the left and centre columns (most notably Subfigures (a) and (b)). The reduced number of fringes indicates a smaller overall deformation. However, the remaining topography is still incompatible with high-finesse operation, requiring further modifications to the membrane geometry.

Change of membrane geometry

Beyond stress compensation, the deformation pattern can be improved by modifying the membrane geometry. Comparing the centre and right columns of Figure 3.3, Subfigures (b) and (c) show two membranes incorporating the high-stress layer that differ only in shape. While both exhibit comparable deflection magnitudes, the circular membrane yields a more symmetric profile with respect to the mirror centre and provides a more favourable starting point for achieving a sufficiently flat mirror area.

This improvement becomes more pronounced when reducing the membrane size. Comparing the first and second rows of Figure 3.3, decreasing the lateral dimensions — edge length for square membranes and diameter for circular membranes (denoted by b) — leads to a noticeably cleaner deflection pattern. A qualitative explanation is that smaller membranes suppress the excitation of higher-order mechanical deformation modes, since

[2]Laseroptik GmbH is a company focusing on optical elements and coatings that is based near Hannover, Germany. Contact: Tobias Groß (tgross@laseroptik.de)

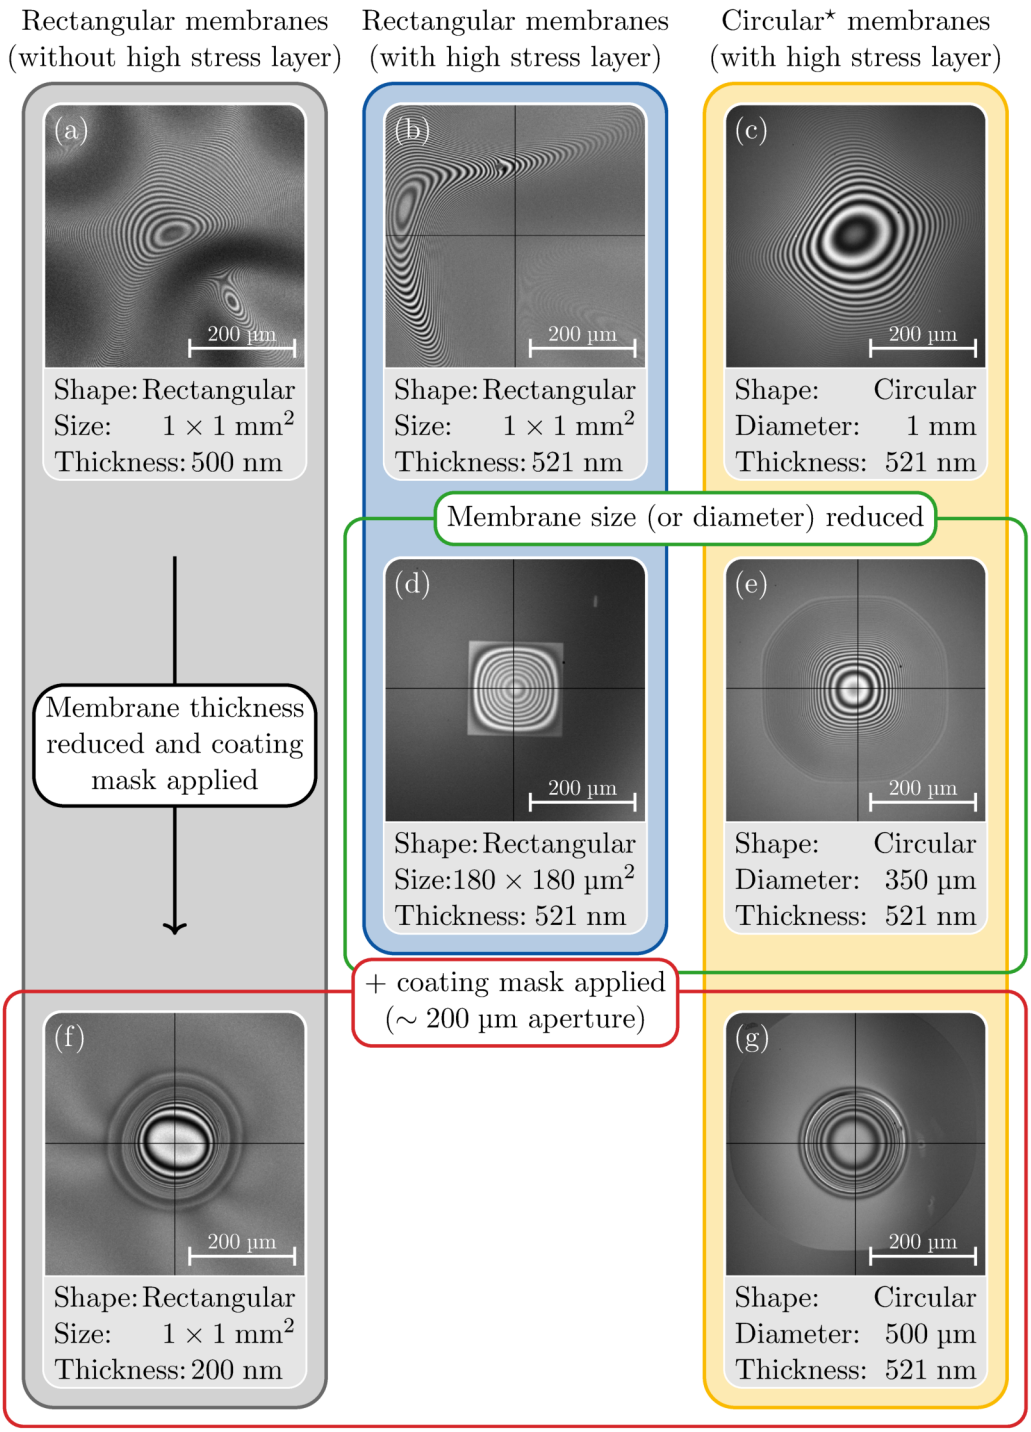


Figure 3.3.: Photographies of the green light interferograms observed on various types of coated silicon nitride membranes, with specifications listed below the individual images. For comparison, the coloured outlines group different membrane characteristics, e.g. reduced size, to show the immediate effects.

*Due to the lattice structure of the silicon wafer the circular shape is a multisided polygon, which becomes apparent towards smaller radii.

the characteristic mode frequencies increase as the size is reduced. For example, the resonance frequency of a mechanical mode scales approximately as [73]:

$$\omega_{\mathrm{M}} \propto b^{-2}, \tag{3.1}$$

such that, for a given amount of stored mechanical energy, fewer higher-order modes contribute to the deformation as b decreases.

In practice, the minimum achievable membrane size, particularly for circular membranes, is limited by the crystalline structure of the silicon frame [66, 74, 75]. For diameters below $b = 350$ µm, the "circular" opening becomes increasingly polygonal. To avoid this limitation while retaining the benefits of reduced size and circular symmetry, the final prototype is based on a circular membrane with a diameter of $b = 500$ µm.

Application of coating mask

While the membrane geometry is constrained by the silicon frame, the coated area can be reduced independently to further limit coating-induced deformation. To achieve this experimentally, we apply a shadow mask prior to the deposition of the high-reflectivity film by placing a thin copper plate[3] with a precisely defined aperture into a dedicated cutout of the coating cartridge. The aperture diameter of 200 µm restricts the mirror coating to a localised region of the membrane. The copper plate thickness of approximately 30 µm ensures minimal shadowing effects during angled deposition, preserving the optical quality of the central coated region.

Interferograms of masked membrane mirrors are shown in the bottom row of Figure 3.3 (highlighted in red). In particular, Subfigure (g) exhibits a central region that appears sufficiently flat for cavity operation. To quantify this, we reconstruct the surface profile using the surface reconstruction algorithm introduced in Section 3.1.1. The resulting profile is shown in Figure 3.4.

To determine the local curvature of the membrane mirror, we fit a parabolic function:

$$\Delta z(r) = -\frac{r^2}{2R_{\mathrm{M}}} + c_1 \cdot r + c_2, \tag{3.2}$$

to the central region of the profile (±10 µm around the centre). Here, R_{M} denotes the radius of curvature, c_1 accounts for residual tilt, and c_2 is an offset. The fits, shown in red on the top and right axes of Figure 3.4, yield radii of curvature of $R_{\mathrm{M,x}} \approx 7.5$ mm and $R_{\mathrm{M,y}} \approx 8.6$ mm, corresponding to an ellipticity of $\epsilon \approx 0.88$.

Finally, we estimate the effective mirror size by evaluating deviations of the measured profile from the fitted paraboloid. Once these deviations exceed $\lambda/4$, the corresponding phase error becomes significant and the affected region no longer contributes efficiently to the cavity mode, providing a practical tolerance for defining the effective mirror size. From the reconstructed surface, this criterion yields an effective mirror diameter of $b_{\mathrm{M}} \approx 96$ µm; even using a more conservative $\lambda/10$ threshold still gives $b_{\mathrm{M}} \approx 95$ µm.

[3] GA200C3 from Ted Pella, Inc.

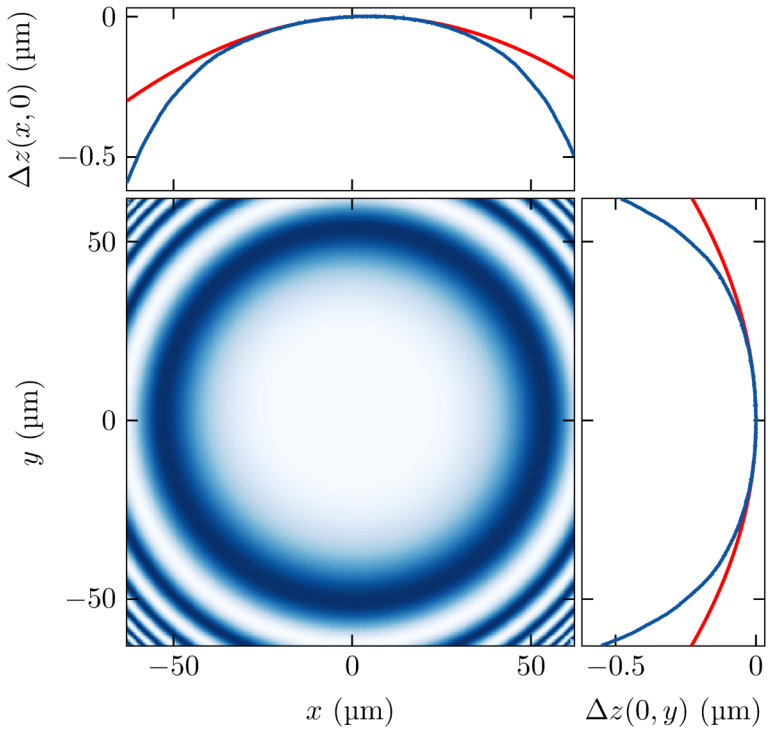


Figure 3.4.: Reconstructed surface profile of the final coated membrane prototype (see Figure 3.3.(g)), obtained from green-light interferometry. Adjacent interference maxima correspond to a height difference of $\Delta z = 532$ nm. The top and left panels show cross-sections through the centre of the reconstructed profile along the x- and y-axes, respectively. Red curves indicate parabolic fits to the central region, which are used to extract the local radii of curvature and to estimate the usable mirror diameter of the coated area.

Final remarks on the membrane mirror

The most pronounced improvement of the mirror surface profile is achieved by restricting the coated area, as illustrated by the prototype depicted in Figure 3.3.(f). Despite using a comparatively thin silicon nitride layer of 200 nm — which, all else being equal, would be expected to increase the susceptibility to deformation — the coated structure exhibits a relatively flat central region. Surface reconstruction yields radii of curvature on the order of $\lesssim 1$ mm. However, reducing the observed ellipticity may further increase the effective radius. The ellipticity is attributed to the lack of radial symmetry in the square membrane geometry and to a slight off-centre placement of the shadow mask, both of which lead to an uneven distribution of coating-induced stress. These observations suggest that in future iterations the compensatory high-stress layer may be omitted, provided that the mask alignment prior to the coating process can be improved.

In contrast, the optimised prototype shown in Figure 3.3.(g) and Figure 3.4 combines all key design elements: a high-stress compensatory layer, a reduced circular membrane geometry, and a masked coating region. This design results in an approximately

parabolic mirror surface with radii of curvature of $R_{\mathrm{M}} \approx 8$ mm. Compared to the fibre mirrors, whose radii of curvature are on the order of a few hundred micrometres (see Table 2.1), the membrane curvature is larger by more than an order of magnitude and therefore has only a minor influence on the cavity stability.

Previous characterisation of the fibre cavity identified clipping losses at larger mirror separations, consistent with an effective fibre mirror radius of $a_1 = 34(1)$ µm (see Section 2.3.2). Here, we find that the membrane mirror provides a usable mirror radius of $a_{\mathrm{M}} = b_{\mathrm{M}}/2 \approx 48$ µm, such that it does not impose a fundamental limitation on the cavity mode size. Overall, these results demonstrate that coating-induced deformations can be sufficiently mitigated and that the realised membrane mirror is suitable for integration into high-finesse cavity experiments.

3.2. Optimisation of the mirror coating composition

The coating applied to the silicon nitride membrane serves as a dielectric Bragg mirror that is composed of alternating high- and low-index layers. At each interface, the incident field is partially reflected and transmitted, and the resulting partial waves interfere. By choosing the layer thicknesses such that each layer has a quarter-wave optical thickness (QWOT, $\lambda_0/(4n)$), constructive interference of the reflected fields is obtained, thereby enhancing the overall reflectivity. Increasing the number of layer pairs then drives the reflectance towards unity, ultimately limited by material absorption. A detailed description of the quarter-wavelength Bragg mirror is provided in Appendix A.1.

To achieve predictable optical performance while optimising for mechanical stability, the coating composition must be tailored to the silicon nitride membrane. In particular, the required reflectivity must be balanced against the total coating thickness, since the coating-induced stress scales with the overall stack height [72]. Because the membrane itself forms part of the multilayer structure, its optical thickness can be incorporated into the design and exploited as an additional layer contributing to the collective interference. To achieve effective contribution, we explore an approach involving a compensatory layer that brings the combined optical thickness to an integer QWOT.

3.2.1. Refractive index of silicon nitride

To optimise a compensatory layer to the silicon nitride relies on precise measurements of the silicon nitride's optical thickness. However, the silicon nitride membranes used in this study are custom prototypes. Therefore neither their refractive index at the target wavelength ($\lambda_0 = 780$ nm) nor their physical thickness were experimentally known, and these values were not specified by the manufacturer. To determine both parameters, we exploit the fact that a single silicon nitride membrane acts as a low-finesse Fabry–Pérot etalon: reflections at the silicon nitride–air interfaces lead to multiple internal reflections and produce a characteristic transmission spectrum.

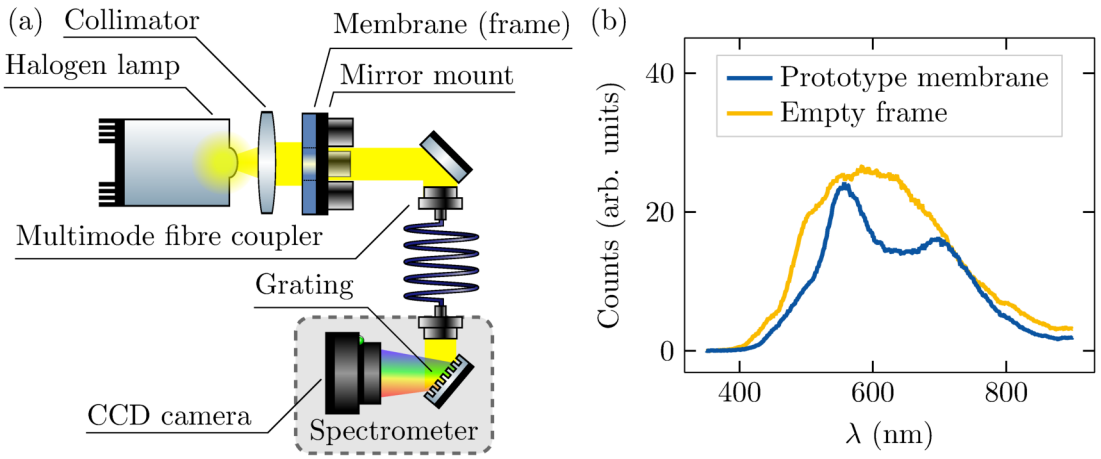


Figure 3.5.: (a) Sketch of the white light transmission spectroscopy setup used for determining the optical thickness of the silicon nitride windows. A halogen lamp is used as white light source and the transmission of the silicon nitride windows is monitored with a fibre-coupled spectrometer. (b) Recorded spectral profile of the halogen lamp, with and without the uncoated prototype membrane in the optical path. This figure includes elements generated with Ref. [54].

Following Equation (2.5b), the transmittance of such an etalon can be written as:

$$\mathcal{T} = \frac{1}{1 + \frac{4\mathcal{R}}{(1-\mathcal{R})^2}\sin^2(k_0 t_{\text{opt}})}, \tag{3.3}$$

where $k_0 = 2\pi/\lambda_0$ is the vacuum wavenumber, $t_{\text{opt}} = nt$ denotes the optical thickness of the silicon nitride layer (with refractive index n and physical thickness t), and $\mathcal{R} = |\rho|^2$ is the Fresnel intensity reflectance at the silicon nitride–air interface (with ρ being the Fresnel amplitude reflection coefficient, as defined in Equation (A.1)). For optical thicknesses on the order of $t_{\text{opt}} \sim 1$ µm [66], the resulting free spectral range is large in wavelength units, requiring transmission spectroscopy over a broad spectral range.

We therefore perform white-light transmission spectroscopy using a halogen light source, as illustrated in Figure 3.5.(a). Light from the lamp is collimated and directed onto the membrane, and the transmission is collected into a multimode fibre and analysed with a commercial spectrometer[4] with a spectral resolution of $\Delta\lambda = 1$ nm. To ensure near-normal incidence, the membrane frame is mounted on a mirror mount and its tilt is adjusted until the light reflected from the silicon frame overlaps with the source.

Because the halogen lamp spectrum is not uniform, we first record a reference spectrum using an empty silicon frame, followed by the transmitted spectrum of the silicon nitride membrane prototype under identical conditions, as compared in Figure 3.5.(b). Normalising the membrane spectrum by the reference eliminates the spectral envelope of the light source and yields the membrane transmittance shown in Figure 3.6.(a).

In Figure 3.6.(a), the red curve shows a fit of Equation (3.3) to the data within the grey-

[4] Red Tide USB650 from Ocean Optics, Inc.

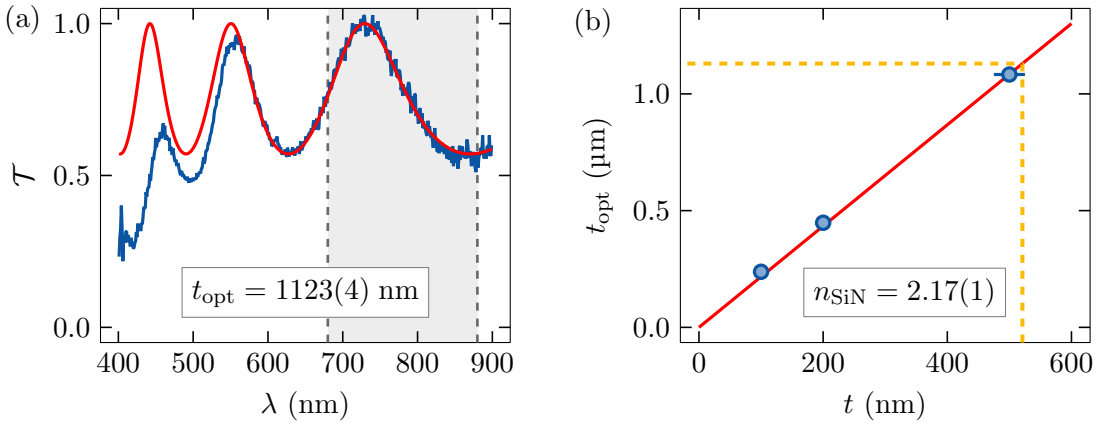


Figure 3.6.: (a) Normalised transmittance spectrum of a prototype silicon nitride membrane window including the high-stress layer. The red curve shows a fit to the etalon model in the range 680 nm – 880 nm (grey shaded region), chosen to minimise the influence of dispersion. (b) Extracted optical thicknesses of silicon nitride membranes with nominal physical thicknesses $t \in \{100\ \mathrm{nm}, 200\ \mathrm{nm}, 500\ \mathrm{nm}\}$. The solid line is a linear fit used to determine the refractive index n_{SiN}. The orange lines indicates the mapping of the optical thickness measured for the prototype in (a) and the corresponding inferred physical thickness.

shaded region. From this fit, we extract an optical thickness of $t_{\mathrm{opt}} = 1123(4)$ nm for the prototype membrane. Data outside the fitting window are excluded because the model neglects dispersion: strictly speaking, both k_0 and t_{opt} depend on wavelength. Focusing on the vicinity of $\lambda_0 \approx 780$ nm allows these effects to be negligible to good approximation, whereas deviations increase at larger detuning. In addition, increased absorption is observed towards the ultraviolet region [76]. Notably, the prototype membrane consists of two silicon nitride layers, while Equation (3.3) assumes a single layer. The excellent agreement within the fit region indicates that the bi-layer membrane can be modelled effectively by a single layer with an effective optical thickness.

To separately determine the refractive index n and physical thickness t, we repeat the measurement for membranes with nominal thicknesses of 100 nm, 200 nm, and 500 nm. The resulting optical thicknesses are shown in Figure 3.6.(b). The horizontal error bars reflect the manufacturer-specified 5 % tolerance on the physical thickness [66], while the vertical uncertainties correspond to the standard deviation of fits on various membranes with the same thickness and are negligible. A linear fit yields a refractive index of $n_{\mathrm{SiN}} = 2.17(1)$, consistent with the manufacturer's specified range of $n_{\mathrm{SiN}} = 2.15$–2.17 at $\lambda_0 = 630$ nm [66]. Using this value, we infer a physical thickness of $t_{\mathrm{SiN}} \approx 521$ nm for the prototype membrane, as indicated by the orange lines in Figure 3.6.(b).

3.2.2. Optimised coating composition

For the membrane prototype investigated here, the measured silicon nitride thickness corresponds to approximately 5.77 QWOT at $\lambda_0 = 780$ nm, which is not optimal for achieving maximal reflectivity in a conventional Bragg design. To compensate this mismatch, we introduce an additional dielectric layer prior to the quarter-wave stack. Together with the silicon nitride membrane, this compensation layer then forms an effective low-finesse etalon. By choosing the combined optical thickness of the Si_3N_4 membrane and the compensation layer to be 7 QWOT, we deliberately detune this etalon from resonance, reducing the optical intensity within the membrane and thereby reducing absorption and potential nonlinear effects in the silicon nitride layer.

To reduce internal reflections at the silicon nitride–compensation layer interface, the difference in refractive index should be as small as possible, making the high-index material of the dielectric coating a natural choice. High-reflectivity dielectric mirrors in the optical range are typically realised using alternating layers of silicon dioxide (SiO_2) as the low-index material and either tantalum pentoxide (Ta_2O_5) or niobium pentoxide (Nb_2O_5) as the high-index material. Both high-index candidates are suitable for the compensation layer, but here we choose Nb_2O_5 due to its comparatively higher refractive index [77]. For a given optical thickness, this also reduces the layer's physical thickness and thus lowers the total stack height, resulting in less coating-induced stress [72].

As a result, the final coating design consists of the 521 nm silicon nitride membrane followed by an Nb_2O_5 compensation layer of approximately 1.22 QWOT, yielding a combined thickness of 7 QWOT. On top of this, nine pairs of alternating SiO_2 and Nb_2O_5 quarter-wave layers are deposited. The stack is terminated by an additional SiO_2 quarter-wave layer, which protects the exposed Nb_2O_5 surface from environmental degradation. The total stack thickness, including the silicon nitride membrane, is $t_M \approx 2700$ nm, i.e. the coating exceeds the membrane thickness by approximately a factor of four. This highlights the substantial mechanical load applied to the membrane. However, as discussed in Section 3.1, the adopted geometry optimisation strategies mitigate its impact on mechanical stability.

Figure 3.7 shows the simulated coating response obtained using the *transfer-matrix method* (see Appendix A), alongside the transmission spectra of an uncoated and a coated membrane, measured using white-light spectroscopy. The simulation agrees closely with the experimental spectra, confirming the successful application of the high-reflectivity coating. The increasing deviation at shorter wavelengths is attributed to dispersion, which is not included in the numerical model.

The inset in Figure 3.7 shows a transmission spectrum recorded using a line filter that transmits only a narrow band around $\lambda = 780$ nm. This suppresses high-intensity spectral components that would otherwise saturate the spectrometer, enabling longer integration times and improved sensitivity near the design wavelength. From this measurement, we extract a transmittance of $\mathcal{T} = 452(47)$ ppm for the coated membrane.

Taken together, these results demonstrate that the refined membrane mirror combines a low-curvature surface profile with high reflectivity, making it suitable for integration into high-finesse cavity experiments.

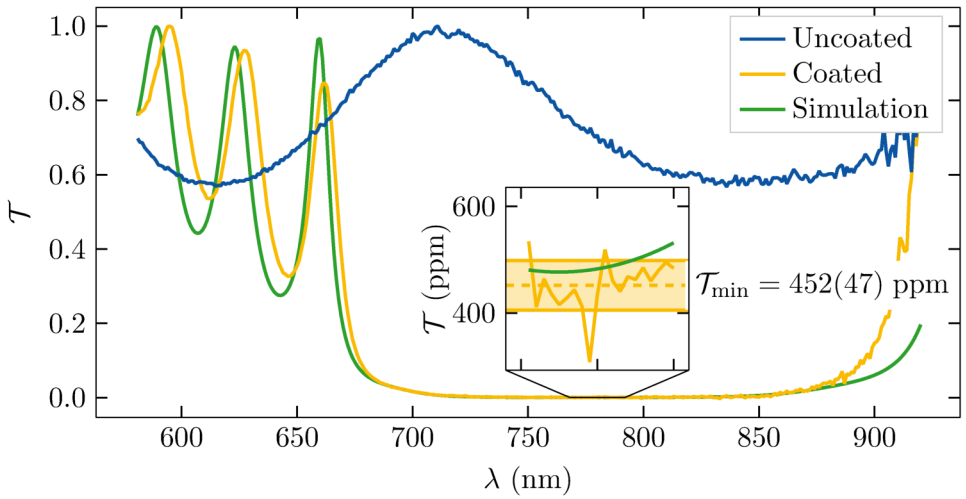


Figure 3.7.: Comparison of the transmission spectra of uncoated (blue) and coated (orange) silicon nitride window (including the high stress layer), alongside, the simulated transmittance (green). The inset corresponds to a measurement with higher integration times of the spectrometer, enabling higher resolution. The shaded area indicates the standard deviation around the average minimal transmittance.

3.3. Optical test cavity with a membrane mirror

The membrane mirror prototypes developed in this chapter exhibit strong potential as coupling mirrors in a coupled-cavity configuration. However, in a high-finesse resonator, even small imperfections can be amplified by the large number of intracavity round trips. In particular, scattering from surface roughness or excess absorption may remain undetected in the single-pass measurements presented above, yet still prevent the formation of a high-finesse cavity. To assess the mirror quality under realistic cavity conditions prior to integration into the final setup, we therefore construct a test cavity between a fibre mirror and the membrane surface. Since the membrane mirror is asymmetric, with the silicon nitride layer present only on one side, both membrane surfaces are characterised independently. To serve as reference for later investigations, each membrane side is analysed using a different fibre mirror, reflecting their positions in the final coupled-cavity configuration.

For the test setup, the silicon frame is gently clamped in a dedicated holder, while the fibre mirror is positioned opposite the membrane using a five-axis translation stage (three translational and two angular degrees of freedom). A piezoelectric actuator integrated into the stage provides fine control of the cavity length. Cavity alignment and linewidth measurements are performed using the same procedure as for the fibre–fibre cavity described in Section 2.3: coarse alignment and an initial estimate of the cavity length are obtained using a USB microscope, while fine scans of the cavity length are calibrated in frequency units via EOM-generated sidebands (see Section 2.3.1).

For measurements on the backside of the membrane, optical access is partially ob-

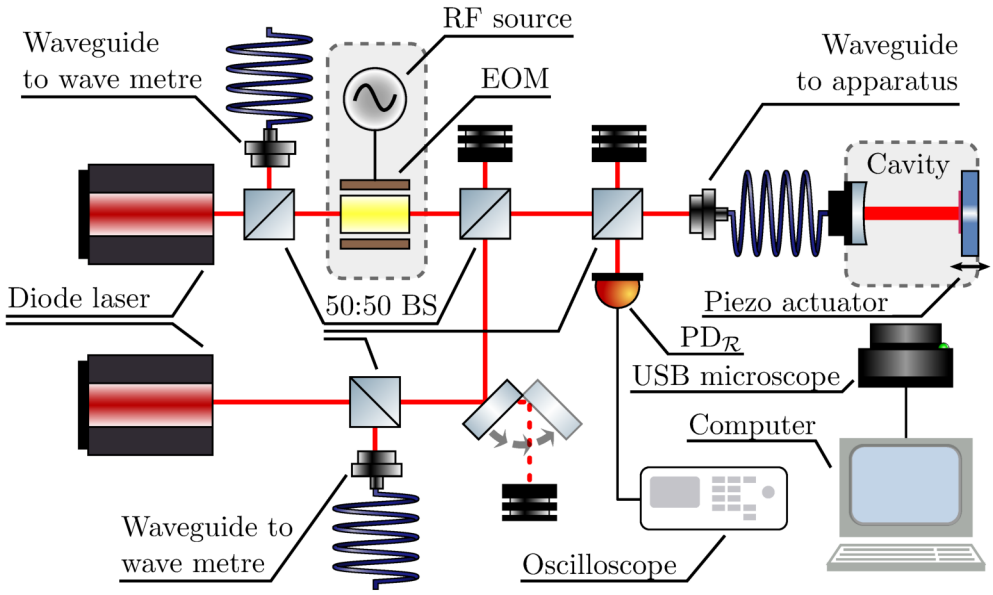


Figure 3.8.: Sketch of the setup used for spectroscopy of the fibre-membrane cavity. The cavity is detuned via the piezoelectric actuator in the mount of the membrane mirror. The scan is calibrated by using sidebands, generated by phase modulation with an EOM. Furthermore, the setup features an USB microscope and a second laser to measure cavity length and free spectral range, respectively. The wavelengths are measured by a commercial wavelength metre with MHz resolution. This figure includes elements generated with Ref. [54].

structed by the silicon frame, which complicates direct length estimation. To determine the free spectral range in this configuration, we optionally combine a second laser with the probe laser. When both lasers are simultaneously resonant with the cavity and the detuning between them is small, so that dispersion effects are negligible, this frequency difference corresponds to an integer multiple of the free spectral range. The absolute laser frequencies are measured using a wavelength meter[5]. A sketch of the extended optical setup is shown in Figure 3.8.

Reflection spectra recorded at a cavity length of ~ 150 µm are shown in Figure 3.9, with the upper and lower subfigures corresponding to cavities formed with the front and back side of the membrane, respectively. The detuning axis is calibrated using EOM sidebands (with modulation frequencies of 1 GHz and 3 GHz in the two measurements). We fit the spectra with Lorentzian line shapes to extract the resonance linewidth $\delta\nu$. Combined with the measured free spectral range ν_{FSR}, the finesse is then obtained from:

$$\mathcal{F} = \frac{\nu_{\mathrm{FSR}}}{\delta\nu}, \tag{3.4}$$

or equivalently $\mathcal{F} = c/(2d\,\delta\nu)$ when the cavity length d can be estimated visually. The extracted parameters are summarised in Table 3.1.

[5] WS6 series from HighFinesse GmbH

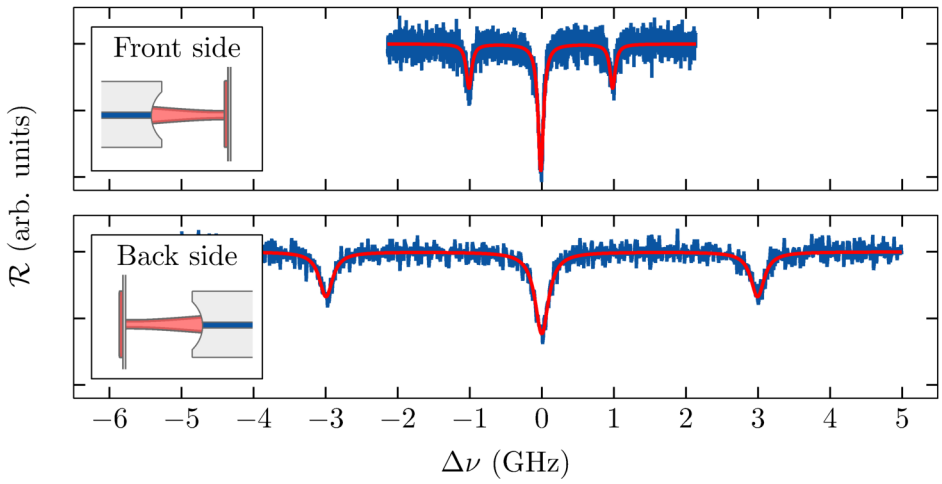


Figure 3.9.: Normalised cavity reflection spectra as a function of the detuning $\Delta\nu$ between the incident laser frequency and the cavity resonance. The red curve shows a fit consisting of three Lorentzian line shapes (carrier and calibration sidebands), from which the resonance linewidth $\delta\nu$ is extracted. Both measurements were performed at a cavity length of approximately $d \sim 150$ µm, but with opposite orientations of the membrane mirror, as indicated by the insets.

Table 3.1.: Summary of characteristic cavity parameters for the two orientations of the membrane mirror. The table lists the fit parameters extracted from the reflection spectra shown in Figure 3.9, together with the average finesse $\mathcal{F}_{\mathrm{Av}}$ obtained from the plateau values in Figure 3.10. The corresponding non-radiative losses $\mathcal{L}$ are inferred from the deviation from the loss-limited finesse expected for mirror transmittances of $\mathcal{T}_1 = 124(11)$ ppm and $\mathcal{T}_2 = 452(47)$ ppm. Quantities determined from direct measurements are marked with *.

		Front side	Back side
ν_{FSR}	(THz)	0.99(3)	0.97(0)*
d	(µm)	152(5)*	154.63(0)
ν_{FWHM}	(MHz)	89(1)*	224.8(6)*
$\mathcal{F}$	-	11003(410)	4315(111)
$\mathcal{F}_{\mathrm{Av}}$	-	11090(697)	4032(435)
$\mathcal{L}$	(ppm)	−1(65)	990(176)

A key observation is that the cavity formed with the backside of the membrane exhibits a substantially broader resonance, indicating increased intracavity losses. To quantify this behaviour, we repeat the measurement for varying cavity lengths and extract the finesse $\mathcal{F}(d)$, shown in Figure 3.10. The backside cavity (orange) displays a lower but constant finesse over the investigated range. In contrast, the finesse of the cavity formed with the front side (blue) decreases rapidly beyond $d \sim 175$ µm. This behaviour is characteristic of additional length-dependent loss channels that arise when the transverse cavity mode expands with increasing mirror separation. In particular, once the mode size

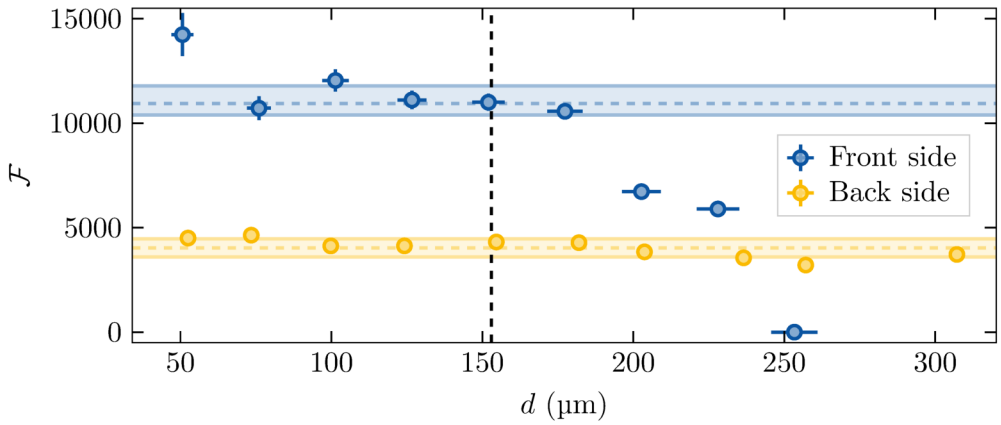


Figure 3.10.: Finesse $\mathcal{F}$ of the cavity for different cavity lengths d with the membrane's front and back side facing the fibre mirror, respectively. The data points close to the vertical dashed line were shown in Figure 3.9, corresponding to a cavity length of $d \approx 150$ µm. The horizontal shaded areas indicate the average finesse of the two formed cavities, neglecting clipping losses on the fibre end-mirrors.

at the fibre mirror becomes comparable to the effective mirror size, the field is partially clipped, introducing excess loss and reducing the finesse, as discussed in Section 2.3.2. Since the membrane mirror provides a larger usable size (see Section 3.1) and does not constrain the mode at these lengths, the observed loss is consistent with clipping being dominated by the fibre mirror. This supports the assumption introduced in Section 2.3.2, where clipping losses were attributed predominantly to only one fibre end-mirror.

Since the finesse is expected to be independent of cavity length in the absence of additional loss channels, we estimate representative values from the plateau regions of the respective datasets. The corresponding mean values and standard deviations are indicated by the shaded regions in Figure 3.10.

3.3.1. Absorption within the silicon nitride layer

The measured finesse values of the two test cavities reveal a pronounced asymmetry in intracavity losses. This behaviour remains unchanged when flipping the membrane or replacing it with another coated sample, ruling out contamination artefacts and indicating a systematic origin. Since the fibre mirrors were characterised independently prior to these measurements (see Section 2.3.2), we attribute the additional losses to the membrane mirror itself.

The two investigated cavity configurations differ only in the orientation of the membrane mirror, i.e. in the layer order encountered by an incident field. This distinction is relevant because a Bragg mirror is not a perfectly reflecting surface: the optical field penetrates into the multilayer stack, resulting in a residual intensity within the coating [41, 78]. Importantly, this residual intensity is not distributed uniformly across the layers, but is largest near the outermost layers and decays rapidly with depth, as demonstrated

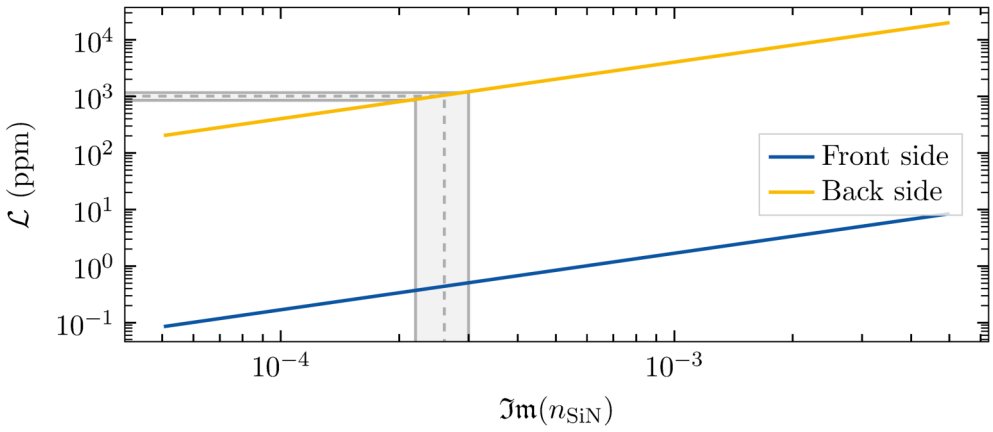


Figure 3.11.: Simulation of the relationship between the imaginary part of the refractive index of silicon nitride $\mathfrak{Im}(n_{\mathrm{SiN}})$ and the (absorptive) losses of the mirror, $\mathcal{L}$, for both considered orientations of the membrane mirror. The basis principles of the numerical analysis are outlined in Appendix A. The shaded area maps the measured absorptive losses to a corresponding imaginary part.

in Appendix A.2.

Consequently, when the silicon nitride layer forms the leading layer of the stack, with respect to the incident light field, the optical intensity within this material is maximised. In the reversed orientation, the silicon nitride layer is located deeper within the structure and experiences a substantially reduced intensity due to the preceding decay of the field. If silicon nitride introduces optical loss, these differences in local intensity translate directly into different effective loss rates for the two cavity configurations.

While this analysis does not distinguish between scattering and absorption, the nonlinear cavity response investigated in Chapter 6 strongly suggests a photothermal effect that is predominantly driven by absorption. We therefore model the excess loss as absorption in the silicon nitride layer by introducing an imaginary component to its refractive index.

Using the *transfer-matrix method* (see Appendix A), we compute the reflectance $\mathcal{R}$ and transmittance $\mathcal{T}$ of the membrane mirror and estimate the corresponding non-radiative loss as:

$$\mathcal{L} = 1 - \mathcal{R} - \mathcal{T}, \tag{3.5}$$

where $\mathcal{L}$ is evaluated as a function of the imaginary part of the silicon nitride refractive index. The resulting simulated loss curves are shown in Figure 3.11, where the shaded region indicates the mapping of the experimentally observed cavity losses.

To extract the total intracavity loss from the measurements, we use the fact that the finesse directly quantifies the cavity lifetime and is therefore closely related to the

round-trip loss. Using Equation (2.9), we obtain:

$$\mathcal{F} = \frac{2\pi}{\mathcal{T}_1 + \mathcal{T}_2 + \mathcal{L}} \quad \Longleftrightarrow \quad \mathcal{L} = \frac{2\pi}{\mathcal{F}} - \mathcal{T}_1 - \mathcal{T}_2, \tag{3.6}$$

where $\mathcal{T}_1$ and $\mathcal{T}_2$ denote the intensity transmittances of the fibre and membrane mirrors, respectively. The average finesse values and the resulting loss estimates for both configurations are also summarised in Table 3.1, using $\mathcal{T}_1 = 124(11)$ ppm (see Section 2.3) and $\mathcal{T}_2 = 452(47)$ ppm (see Section 3.2.2).

For the cavity formed with the backside of the membrane, the measured losses correspond to an imaginary refractive index of $\mathfrak{Im}(n_{\mathrm{SiN}}) = 2.6(4) \times 10^{-4}$, which yields an absorption coefficient of $\alpha_{\mathrm{SiN}} = k_0\, \mathfrak{Im}(n_{\mathrm{SiN}}) = 21(3)\ \mathrm{cm}^{-1}$. Previous studies report slightly lower values of $\mathfrak{Im}(n_{\mathrm{SiN}})$, although at $\lambda = 1550$ nm [64]. Given the strong absorption of silicon nitride towards the ultraviolet [76], reduced absorption at longer wavelengths is expected, which supports the plausibility of this discrepancy. In the opposite configuration, the model predicts losses below 1 ppm, i.e. below the experimental resolution, consistent with the absence of detectable excess losses within the measurement uncertainty.

Despite the presence of absorption in the silicon nitride layer, the membrane mirror supports stable cavity operation with high finesse, demonstrating its suitability as a coupling mirror in a coupled-cavity experiment. Moreover, the absorption-induced losses provide a mechanism for investigating nonlinear dynamics in the coupled-cavity system, which will be the focus of Chapter 6.

3.4. Investigation of the membrane's vibrational modes

The experimental configuration studied here, a silicon nitride membrane placed inside an optical cavity, is inspired by optomechanics experiments [21]. The field of optomechanics investigates the interaction between light and mechanical motion through the coupling of optical and mechanical resonators. By exploiting radiation-pressure forces, mechanical systems can be driven, cooled, and measured with high precision, enabling studies that bridge quantum optics and classical mechanics [25, 30, 61, 79].

In particular, the "membrane-in-the-middle" configuration has become a widely used optomechanical architecture because it combines strong optical confinement with a mechanically compliant element that couples dispersively to the intracavity field [21, 64, 80]. Unlike end-mirror implementations, the membrane position within the standing wave can be chosen to tune the coupling strength and even access regimes where the optomechanical interaction becomes predominantly quadratic in displacement. This flexibility makes membrane-based systems a versatile platform for exploring dynamical backaction effects such as optomechanical cooling and induced transparency [62, 81].

The central component of such systems is the mechanical membrane oscillator, which can take various forms, including vibrating silicon nitride membranes [65] or, as here, a coated membrane mirror. Its resonant frequencies, mechanical quality factors, and mode shapes are key parameters that determine the system performance and are governed by

the oscillator's material properties, geometry, and boundary conditions [73, 82].

The vibrational spectrum of the membrane mirror can be measured using the cavity resonator setup introduced in the previous section. Mechanical motion modulates the effective cavity length and therefore the optical response, leading to measurable intensity fluctuations in the reflected or transmitted light field. By recording these fluctuations with a photodetector and analysing their spectral content, the mechanical resonances of the membrane can be identified.

Based on experimental data of the membrane's vibrational modes, we develop a simplified mechanical model that treats the membrane mirror as two uncoupled thin plates, capturing the dominant resonance modes. This model provides guidance for tailoring the membrane geometry to optimise its mechanical resonance spectrum in the future.

3.4.1. Detection of the membrane motion

The detection of membrane motion relies on resolving small variations in the cavity length, Δd. Direct intensity detection at resonance is inefficient, since the slope of the resonance vanishes at this point and the response is only second order in small displacements, $\delta I(d+\Delta d) \approx 0+\mathcal{O}(\Delta d^2)$. Moreover, due to the symmetry of the resonance lineshape, the sign of a displacement cannot be inferred from the transmitted or reflected intensity. A simple alternative is off-resonant detection on the side of the resonance, where the intensity varies quasi-linearly with detuning. However, this approach remains sensitive to laser power fluctuations, which can mimic cavity-length changes and limit the measurement sensitivity.

A more robust scheme exploits the phase response of the reflected field. Cavity resonance occurs when the round-trip phase accumulation equals an integer multiple of 2π, such that the reflected phase exhibits a zero-crossing at resonance (up to an overall phase defined modulo 2π). This feature provides an intensity-independent reference and suppresses sensitivity to laser power noise. In addition, small detuning from resonance translate into a phase shift that is effectively enhanced by the multiple round-trips of the intracavity field, yielding a highly sensitive probe of Δd. A comparison of the reflected intensity and reflected phase as functions of cavity length is shown in Figure 3.12.(a), where the shaded regions illustrate the noticeably different response to the same cavity-length modulation.

Phase-sensitive detection can be implemented using the Pound–Drever–Hall (PDH) technique [83, 84]. In this method, the incident light field is phase-modulated, generating sidebands around the carrier frequency. The interference of carrier and sidebands upon reflection from the cavity is demodulated at the modulation frequency, yielding a dispersive error signal that is proportional to the cavity detuning. A schematic of our implementation is shown in Figure 3.12.(b), where phase modulation is realised via modulation of the laser current [44], with a frequency of $\Omega \approx 60$ MHz. As a result, the sidebands and the carrier overlap within a cavity linewidth ($\nu_{\mathrm{FWHM}} \gtrsim 100$ MHz) manifesting as a single dispersive feature, as shown in Figure 3.12.(c).

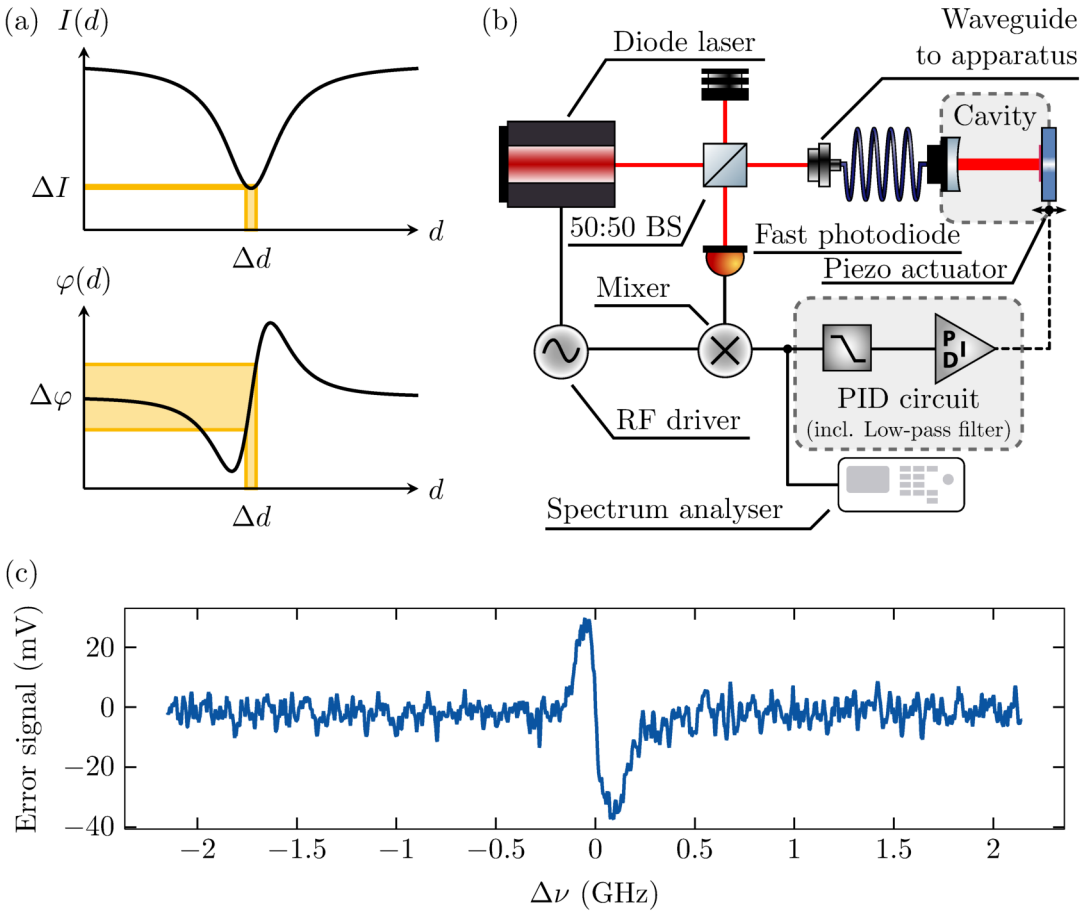


Figure 3.12.: (a) Comparison of the response of the (*top*) reflected intensity $I(d)$ and (*bottom*) acquired phase of the reflected light, $\varphi(d)$, for a given change Δd in the cavity length. (b) Sketch of the implemented PDH setup, facilitating the phase-sensitive measurement. Phase modulation is achieved by modulating the laser current. (c) Generated error signal using a current modulation of $\Omega \approx 60$ MHz, picked off after the PID controller as a function of cavity detuning $\Delta\nu$. The signal exhibits the characteristic symmetry across the zero-crossing. This figure includes elements generated with Ref. [54].

Noise spectrum

The demodulated phase signal (error signal) shown in Figure 3.12.(c) provides a sensitive probe of the membrane motion through its conversion of cavity-length fluctuations into voltage fluctuations. To suppress slow drifts of the cavity resonance, we stabilise the cavity length using a proportional–integral–derivative (PID) controller. The feedback loop has a bandwidth of approximately 2 kHz, such that slow fluctuations are compensated while the mechanical resonances of interest at higher frequencies remain observable [21]. At the same time, locking the cavity enables long acquisition times, which improves the frequency resolution of the spectrum analyser ($\Delta\nu \propto 1/\Delta t$). To record the vibration spectrum while maintaining stabilisation, the error signal is split prior to the PID input,

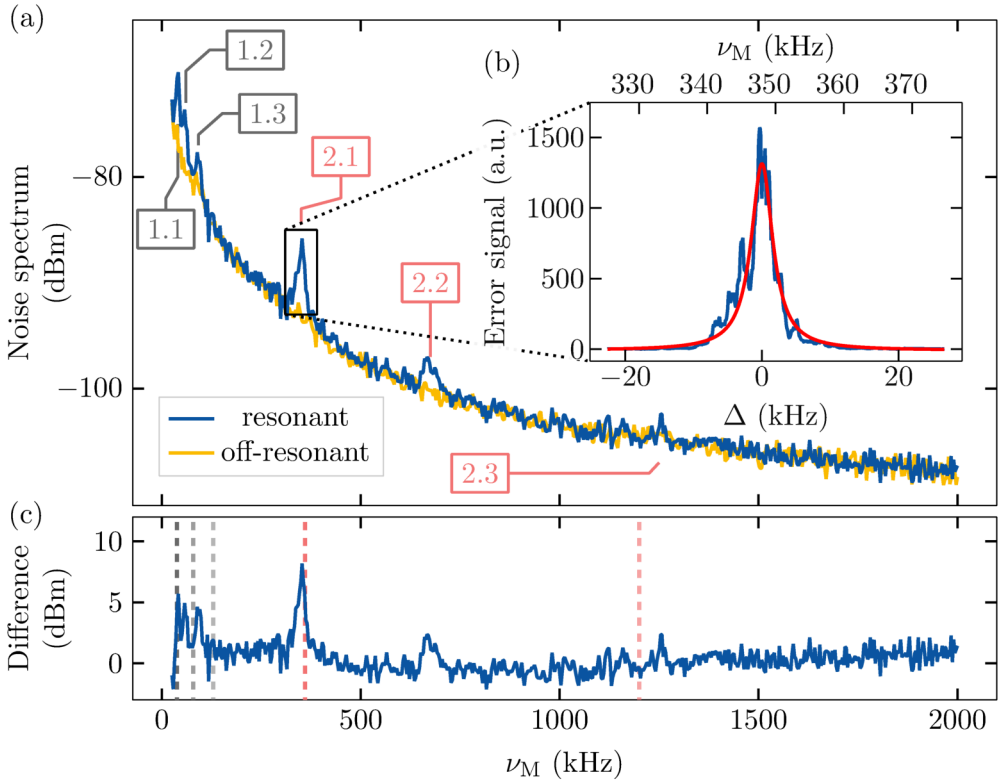


Figure 3.13.: (a) Logarithmic noise spectrum of the cavity reflection signal. To suppress technical contributions unrelated to the cavity, a reference spectrum is recorded with the cavity detuned (off-resonant), and the difference between the locked and reference traces reveals several mechanical resonances. Two distinct mode families are highlighted in different colours. (b) Magnified, high-resolution spectrum of the strongest resonance (labelled 2.1) after background subtraction and displayed on a linear scale. The red curve shows a Lorentzian fit. (c) Background-corrected noise spectrum obtained by fitting a $1/\nu_M$ dependence to the off-resonant trace and subtracting it from the locked measurement. Vertical lines indicate the resonance frequencies predicted by the mechanical model discussed in Section 3.4.2: grey lines correspond to modes of the suspended silicon nitride membrane, and red lines to modes of the coated membrane-mirror region. Higher-order modes are shown with decreasing opacity.

as depicted in Figure 3.12.(b): one branch is sent to a spectrum analyser, while the other is used for feedback control of the cavity length.

The measured noise spectrum is shown in Figure 3.13.(a). To distinguish mechanical resonances from technical background noise (e.g., electromagnetic pickup or electronic resonances), a reference spectrum is recorded with the cavity detuned and unlocked (orange trace). Both spectra exhibit an approximate $1/\nu_M$ dependence, consistent with typical electronic noise behaviour [85]. Importantly, several distinct peaks appear only in the locked configuration, indicating the presence of discrete mechanical modes of the membrane mirror.

As an illustrative example, the inset of Figure 3.13 (Subfigure (b)) shows a high-resolution spectrum around the strongest resonance (labelled 2.1), where the difference between locked and unlocked traces is plotted on a linear scale. Fitting the resonance with a Lorentzian lineshape (red curve) yields a centre frequency $\omega_{\mathrm{M}} = 2\pi \times 347.99(3)$ kHz and a linewidth $\Gamma_{\mathrm{M}} = 2\pi \times 3.84(8)$ kHz. The resulting mechanical quality factor, defined as $Q = \omega_{\mathrm{M}}/\Gamma_{\mathrm{M}} = 91(2)$, is significantly lower than values reported for comparable silicon nitride membrane oscillators [21, 24, 25, 65, 80], despite similar order of magnitude of the resonance frequencies.

This reduced quality factor indicates enhanced damping of the membrane motion, which we attribute primarily to operation at ambient pressure. In contrast to many optomechanical experiments performed under vacuum, the measurements presented here were conducted in air, where viscous damping and molecular collisions introduce additional dissipation. Previous studies have shown that operating thin membrane oscillators in vacuum can increase their quality factor by several orders of magnitude [86]. Accordingly, substantially higher quality factors are expected for the membrane mirror in future measurements under vacuum conditions.

3.4.2. Mechanical model

A characteristic feature of a mechanical oscillator is a family of resonant modes that appear at approximately regular frequency intervals, with higher-order modes forming a comb-like structure relative to the fundamental resonance. In the measured spectrum shown in Figure 3.13, six pronounced peaks follow this pattern: the first three and the last three resonances form two distinct mode families, highlighted by different colours. This observation suggests that the membrane mirror supports two dominant mechanical oscillators with different fundamental frequencies.

The geometry of the membrane mirror motivates a simple description in terms of two mechanical elements: the suspended silicon nitride membrane within the surrounding frame and the coated central Bragg mirror region. As illustrated in Figure 3.14, we describe these contributions by two displacement fields, $u(r,t)$ and $v(r,t)$, corresponding to the motion of the suspended membrane and the flexural motion of the Bragg mirror region, respectively. Schematic representations of the associated vibrational modes are shown in Figure 3.14.(b) and Figure 3.14.(c).

A full mode analysis of such composite structures generally requires numerical modelling [87]. However, for thin plates of simple geometry, in particular, circular plates of thickness t and radius a with $a \gg t$, semi-analytical expressions for the resonance frequencies can be obtained within classical plate theory [73, 82, 88, 89]. For a given vibrational mode, the angular frequency is:

$$\omega_{\mathrm{M}} = \left(\frac{\mu}{a}\right)^2 \sqrt{\frac{D}{2\varrho t}}, \tag{3.7}$$

where ϱ is the material density and μ is a dimensionless eigenvalue determined by the

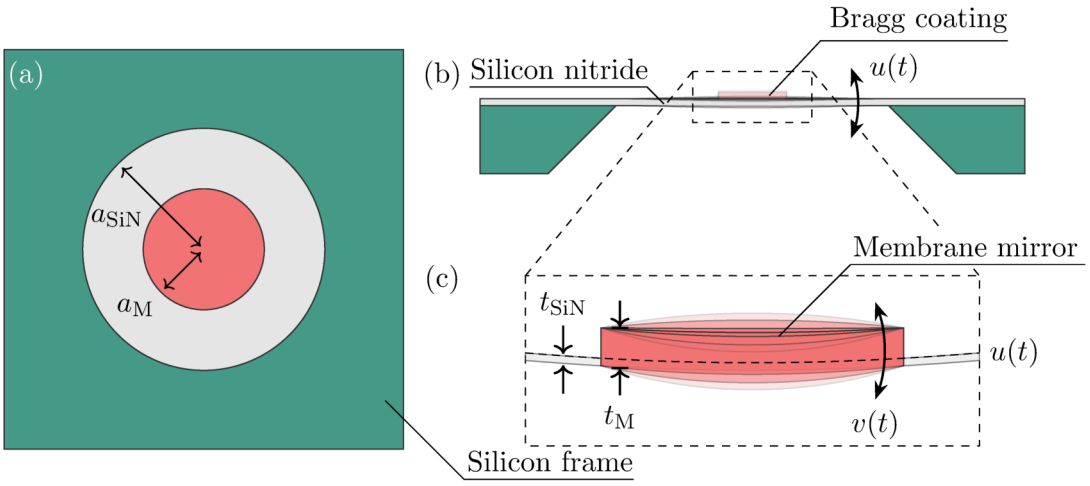


Figure 3.14.: (a) Top and (b) side view of the membrane mirror assembly, consisting of a silicon frame supporting a suspended (partially coated) silicon nitride membrane. The side view illustrates the displacement $u(t)$ associated with the first mechanical oscillator, corresponding to the motion of the suspended silicon nitride. (c) Magnified view of the coated central region, highlighting the displacement $v(t)$ of the second oscillator, which describes the flexural motion of the Bragg-coated mirror area.

plate boundary conditions[6]. Tabulated values of μ for common boundary conditions are available in Ref. [73]. The flexural rigidity is given by:

$$D = \frac{Et^3}{12(1-\nu^2)}, \tag{3.8}$$

where E is Young's modulus, which characterises the stress–strain relationship of the material, and ν is Poisson's ratio, which describes the relationship between axial and transverse strain.

Classical plate theory assumes an isotropic material [88]. This approximation is appropriate for the silicon nitride membrane, whereas the Bragg mirror region is a multilayer composite with spatially varying density and elastic properties. Nevertheless, due to its layered structure, the coating can be approximated as effectively isotropic in the radial plane of the mirror, allowing its mechanical response to be described by homogenised material parameters.

We therefore model the Bragg mirror region using effective material constants obtained from thickness-weighted averages over the total coating thickness t_{M} (see Section 3.2.2).

[6] In the literature, this parameter is commonly denoted as λ. To avoid confusion with the optical wavelength, we use μ throughout this work.

For instance, the effective density can be written as:

$$\begin{aligned}\varrho_{\mathrm{M}} &= \frac{1}{t_{\mathrm{M}}} \sum_j t_j \, \varrho_j \\ &= \frac{\lambda}{4 t_{\mathrm{M}}} \left(\frac{5.78}{n_{\mathrm{SiN}}} \cdot \varrho_{\mathrm{SiN}} + \frac{10.22}{n_{\mathrm{NbO}}} \cdot \varrho_{\mathrm{NbO}} + \frac{10.00}{n_{\mathrm{SiO}}} \cdot \varrho_{\mathrm{SiO}} \right),\end{aligned} \tag{3.9}$$

where t_j and ϱ_j denote the physical thickness and density of each material layer in the stack, respectively. The individual layer thicknesses follow directly from the coating design in units of QWOT at $\lambda_0 = 780$ nm. Using the same procedure, effective values for Young's modulus and Poisson's ratio are estimated. All material parameters used in the model are summarised in Table 3.2.

With these effective parameters, both oscillators can be approximated as circular isotropic plates with different radii, thicknesses, and boundary conditions. For the relevant geometries, the eigenvalues of the fundamental mode satisfy $\mu_0^2 \approx 4.2 - 10.2$, where larger values correspond to more constrained boundary conditions [73]. The resulting frequency ranges for the fundamental modes of both oscillators, together with the underlying material and geometric parameters, are summarised in Table 3.3.

Table 3.2.: Mechanical material parameters of silicon nitride (Si_3N_4), silicon dioxide (SiO_2), and niobium pentoxide (Nb_2O_5), which constitute the dielectric multilayer stack of the membrane mirror.

		Silicon nitride (Si_3N_4)	Silicon dioxide (SiO_2)	Niobium pentaoxid (Nb_2O_5)
E	(GPa)	$270 - 290$ [66, 90]	$72 - 87$ [91]	$130 - 161$ [91]
ϱ	(gcm^{-3})	3.1 [66, 90]	1.98 [91]	4.5 [91]
ν	–	0.20 [90]	$0.11 - 0.17$ [91]	$0.22 - 0.23$ [91]

Table 3.3.: Estimated frequency ranges for the fundamental mechanical resonances of the suspended silicon nitride membrane (mode 1.1) and the coated mirror region (mode 2.1), based on the thin-plate model parameters discussed in this chapter.

		1.1	2.1	
a	(µm)	250	100	
t	(µm)	0.5	2.7	
E	(GPa)	280	140	
ν	–	0.20	0.18	
ϱ	(gcm^{-3})	3.1	3.0	
D	(Nnm)	3	237	
$\frac{\omega_0}{2\pi}$	(kHz)	16	362	$\mu_0^2 = 4.2$
		38	879	$\mu_0^2 = 10.2$

Discussion of the model

A comparison between the analytical model and the experimentally observed fundamental resonances in Figure 3.13 shows good agreement. In particular, the upper bound of the predicted frequency range for the first oscillator and the lower bound for the second coincide well with the measured fundamental resonances. In the model, these bounds are set by the effective stiffness of the surrounding structure, i.e. by the boundary conditions imposed on each plate-like oscillator.

As summarised in Table 3.3, the bending stiffness D of the silicon nitride membrane and that of the coated membrane mirror differ by approximately two orders of magnitude, such that the membrane mirror experiences softer boundary conditions at the interface and correspondingly smaller eigenvalues μ^2. In contrast, the silicon nitride membrane is embedded in a solid silicon frame that is roughly three orders of magnitude thicker, which effectively clamps the membrane and results in stiffer boundary conditions, i.e. larger values of μ^2. The measured resonance frequencies are therefore consistent with the mechanical hierarchy set by the system geometry ($\omega_{\mathrm{M}} \propto \mu^2$, see in Equation (3.7)).

To enable a direct comparison of higher-order modes, Figure 3.13.(c) shows a spectrum, that was background-corrected by subtracting a fitted $1/\nu_{\mathrm{M}}$ contribution from the locked measurement. The resulting peaks are compared to the predicted resonance frequencies, indicated by vertical lines (grey for membrane modes and red for membrane-mirror modes, with higher-order modes shown at lower opacity). While the predicted fundamental resonances agree well with the data, the model systematically overestimates higher-order mode frequencies, indicating that the simplified plate description does not capture the full anharmonic mode structure.

Several effects may contribute to these deviations, including the slightly off-centre mirror position within the frame (see Figure 3.3.(g)), which breaks the rotational symmetry assumed in the model, and mechanical coupling between the suspended membrane region and the coated mirror region, which can hybridise modes. A quantitative treatment would require numerical modelling of the full composite structure. Nevertheless, the present analytical model already provides valuable physical insight.

For instance, low-frequency resonances can limit the observability of optomechanical effects because they are more susceptible to environmental noise and decoherence. According to the model, such modes could be suppressed by fully coating the membrane, which eliminates the mechanically distinct uncoated region and therefore removes the corresponding low-frequency oscillator from the spectrum. However, restricting the coating area reduces stress-induced deformation and improves optical surface quality as discussed in Section 3.1. An application-dependent trade-off therefore exists between boundary stiffness (suppressing low-frequency modes) and coating-induced stress (governing mirror curvature and optical performance).

3.5. Summary

In this chapter, we developed and characterised a highly reflective membrane mirror prototype, achieving a reflectance of $\mathcal{R} \gtrsim 99.87(2)$ %. This was realised by depositing an optimised dielectric Bragg stack onto commercially available silicon nitride membranes. The coating design consists of an Nb_2O_5 compensation layer followed by alternating $\lambda/4$ layers of SiO_2 and Nb_2O_5. White-light transmission spectroscopy confirms the high optical performance of the mirror, yielding a low transmittance of $\mathcal{T} = 452(47)$ ppm.

To mitigate stress-induced deformations introduced by the multilayer coating, we investigated several strategies to improve the surface profile of the membrane mirror. While introducing a high-stress compensation layer resulted only in minor improvements, modifying the membrane geometry — in particular reducing its size and adopting a circular shape — significantly enhanced surface uniformity. Further refinement was achieved by masking the membrane during coating, thereby restricting the coated area. Together, these measures yielded a membrane mirror with an effective parabolic mirror diameter of $b_\mathrm{M} \approx 96$ µm, a radius of curvature of $R_\mathrm{M} \approx 8$ mm, and an ellipticity of $\epsilon \approx 0.88$.

The mirror performance was further evaluated in a cavity configuration formed between the coated membrane and the end facet of a fibre mirror. When the coated surface faced the fibre mirror, the cavity operated in good agreement with theoretical expectations. In the opposite orientation, however, moderate losses of $\mathcal{L} = 990(176)$ ppm were observed.

The observed loss asymmetry was modelled as optical absorption in the silicon nitride layer. By mapping the measured cavity losses onto a *transfer-matrix* model and combining this with the independently determined optical thickness, we estimate the complex refractive index of silicon nitride as $n_\mathrm{SiN} = 2.17(1) + i\,2.6(4) \times 10^{-4}$. Considering the different wavelengths used in previous studies, this value is in reasonable agreement with literature [64, 66].

Beyond optical characterisation, the cavity setup was used to investigate the vibrational modes of the membrane mirror. The measured noise spectrum indicates the presence of two dominant mechanical oscillators, which can be identified with the suspended silicon nitride membrane and the coated mirror region. The fundamental resonance frequencies were found to be $\omega_\mathrm{SiN} \approx 2\pi \times 40.5$ kHz and $\omega_\mathrm{M} = 2\pi \times 347.99(3)$ kHz. For the membrane-mirror mode, we measure a linewidth of $\Gamma_\mathrm{M} = 2\pi \times 3.9(1)$ kHz, corresponding to a mechanical quality factor of $Q = 91(2)$.

While reported quality factors in comparable optomechanical experiments are typically several orders of magnitude higher [21, 64, 80], operation in vacuum has been shown to yield substantial improvements for thin membrane oscillators [86]. Accordingly, future measurements under vacuum conditions are expected to significantly enhance the mechanical performance of the membrane mirror and broaden its applicability in optomechanical coupled-cavity experiments.

4. Realisation of two optical cavity resonators coupled via shared end mirror

Contents

4.1. Mathematical description of two coupled resonators 52
4.1.1. Mode splitting 55
4.1.2. Estimation of the coupling strength 56
4.2. Experimental implementation of the mirror-in-the-middle setup 57
4.2.1. Final sub-cavity performance 62
4.3. Spectroscopy and characterisation near degeneracy point . . 63
4.3.1. Scheme of measurement 64
4.3.2. Measurement of the avoided crossing 68
4.4. Summary and future improvements 70

The previous two chapters characterised the fibre-based microcavity and the coated silicon nitride membrane mirror, respectively. In this chapter, we combine these building blocks by inserting the membrane mirror into the fibre-based resonator, thereby realising a coupled-cavity configuration in which the optical modes of the two sub-cavities are coupled via transmission through the shared membrane. Near degeneracy, the interaction between the cavity resonances gives rise to hybridised modes that extend across both resonators, providing the basis for a photonic-molecule description of the system as a purely optical analogue of light–matter interaction [15, 16].

The physical dimensions of the membrane mirror, in particular its usable mirror size and radius of curvature, impose stringent constraints on the outer cavity geometry. To ensure cavity stability (see Section 2.2.2), short cavity lengths and strongly curved end mirrors are required, which lie within the typical operating range of fibre-based microcavities. However, the resulting cavity length scale becomes comparable to the thickness of the silicon frame that supports the coated membrane, making membrane integration challenging while maintaining precise alignment of the fibre mirrors.

To address this constraint, we employ a passively aligned cavity architecture based on glass ferrules, which enables the fibres to be retracted well beyond the stable operating length and thus provides the clearance required for safe membrane insertion [21]. We first establish a theoretical framework for the coupled-cavity system, yielding quantitative predictions for the expected optical response and forming the basis for subsequent analysis. We then describe the assembly of the apparatus and introduce a measurement

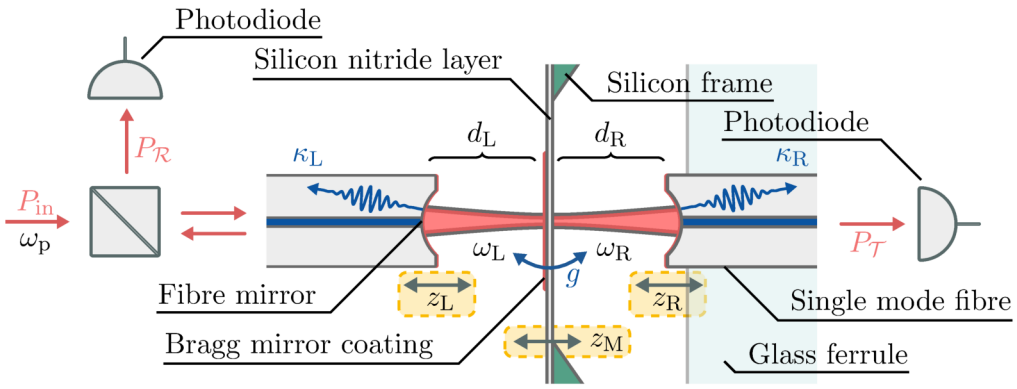


Figure 4.1.: Schematic overview of the experimental apparatus, demonstrating the fibre-based cavity split into two equal sub-cavities using a coated silicon nitride membrane. All characterising quantities are figuratively defined in the image.

scheme designed to acquire a high-resolution, two-dimensional spectrum. Finally, we present experimental evidence of coherent cavity–cavity coupling by resolving and characterising an avoided crossing in the measured response, a hallmark of strongly coupled dissipative systems [6, 92].

4.1. Mathematical description of two coupled resonators

The insertion of a membrane mirror in the centre of a microcavity creates two coupled cavity resonators, one on each side of the coupling element. These cavities are characterised by their respective resonance frequencies, ω_{L} and ω_{R}, and linewidths, κ_{L} and κ_{R}. Photon transmission through the shared membrane mirror enables coupling between the sub-cavities at a rate g. A schematic representation of our coupled-cavity implementation, employing a fibre-based microcavity and a coated silicon nitride membrane, is shown in Figure 4.1.

Mathematically, a single confined optical mode within a cavity resonator is analogous to a harmonic oscillator [93], allowing its Hamiltonian to be expressed as:

$$\hat{\mathcal{H}} = \hbar\omega \cdot \left(\hat{a}^{\dagger}\hat{a} + \frac{1}{2}\right), \tag{4.1}$$

where $\hbar\omega$ represents the energy of each photon in the mode, and the total energy of the mode is given by the product of the photon number and the energy per photon. The photon number is encoded in the number operator $\hat{a}^{\dagger}\hat{a}$, while $\hat{a}^{\dagger}$ and $\hat{a}$ are the creation and annihilation operators, respectively, which add or remove a photon from the corresponding mode [94]. The term $\hbar\omega/2$ corresponds to the zero-point energy, a direct consequence of the Heisenberg uncertainty principle [94]. In the following, we omit this constant offset, as it shifts all energy levels by the same amount and thus generates only an overall global phase in the evolution. Consequently, it does not affect dynamics

in the semi-classical description employed here.

In this case, for a system of two coupled cavities of approximately equal length, with a coupling rate g, the Hamiltonian can be expressed as:

$$\begin{aligned}\hat{\mathcal{H}} = &\overbrace{\hbar\omega_{\mathrm{L}}\left(\hat{a}_{\mathrm{L}}^{\dagger}\hat{a}_{\mathrm{L}}\right) + \hbar\omega_{\mathrm{R}}\left(\hat{a}_{\mathrm{R}}^{\dagger}\hat{a}_{\mathrm{R}}\right)}^{\text{Left- and right-cavity terms}} + \overbrace{\hbar g\left(\hat{a}_{\mathrm{L}}^{\dagger}\hat{a}_{\mathrm{R}} + \hat{a}_{\mathrm{L}}\hat{a}_{\mathrm{R}}^{\dagger}\right)}^{\text{Coupling term}} \\ &+ \underbrace{i\hbar\left(F\mathrm{e}^{i\omega_{\mathrm{p}}t}\hat{a}_{\mathrm{L}} + F^{\star}\mathrm{e}^{-i\omega_{\mathrm{p}}t}\hat{a}_{\mathrm{L}}^{\dagger}\right)}_{\text{Pump term}},\end{aligned} \tag{4.2}$$

where the coupling term $\hbar g\left(\hat{a}_{\mathrm{L}}^{\dagger}\hat{a}_{\mathrm{R}} + \hat{a}_{\mathrm{L}}\hat{a}_{\mathrm{R}}^{\dagger}\right)$ describes the exchange of photons between the left (L) and right (R) cavity modes, and the pump term represents interaction of the left-cavity mode with an incident pump field of frequency ω_{p} and amplitude F [95].

The field intensity, or population, in both cavity modes is represented by the expectation values of the operators $\hat{a}_{\mathrm{L}}$ and $\hat{a}_{\mathrm{R}}$, respectively. To effectively describe the system's dynamics, it is beneficial to describe the system in the Heisenberg picture, where time dependence is attributed to the operators themselves rather than to the state vectors, as in the standard Schrödinger picture [96]. In the Heisenberg framework, the time evolution of the operators is then governed by the general Heisenberg equation of motion:

$$\frac{\mathrm{d}}{\mathrm{d}t}\hat{a}_{\mathrm{L/R}} = \frac{i}{\hbar}\left[\hat{\mathcal{H}}, \hat{a}_{\mathrm{L/R}}\right] + \left(\frac{\partial}{\partial t}\hat{a}_{\mathrm{L/R}}\right), \tag{4.3}$$

where, $\hat{\mathcal{H}}$ is the above Hamiltonian, and $[\cdot,\cdot]$ denotes the commutator, emphasising the non-commutative properties of the operators. The second term allows us to model dissipative effects [94]. In Section 2.1, we showed that the cavity resonance is well described by a Lorentzian lineshape in the frequency domain, which corresponds to an exponentially decaying intracavity field in the time domain. Phenomenologically, this damping can be captured by adding a term proportional to the field operator, such that $\partial_t\hat{a}_{\mathrm{L/R}} \to (\kappa_{\mathrm{L/R}}/2)\hat{a}_{\mathrm{L/R}}$.

The Heisenberg picture is particularly advantageous in the semi-classical limit, where operators can be replaced by their expectation values, $\hat{a}_{\mathrm{L/R}} \to \langle\hat{a}_{\mathrm{L/R}}\rangle \equiv \alpha_{\mathrm{L/R}}$, bypassing the need for explicit wavefunction calculations [97]. Consequently, the time evolution of the cavity fields, following from the evaluated Heisenberg equations of motion, yields the coupled differential equations:

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix}\alpha_{\mathrm{L}}\\ \alpha_{\mathrm{R}}\end{pmatrix} = \underbrace{\begin{pmatrix} i(\Delta_{\mathrm{L}} - \Delta_{\mathrm{p}}) - \frac{\kappa_{\mathrm{L}}}{2} & -ig \\ -ig & -i\Delta_{\mathrm{p}} - \frac{\kappa_{\mathrm{R}}}{2}\end{pmatrix}}_{\tilde{\mathcal{H}}} \cdot \begin{pmatrix}\alpha_{\mathrm{L}}\\ \alpha_{\mathrm{R}}\end{pmatrix} + \underbrace{\frac{\kappa_{\mathrm{L}}}{2}\begin{pmatrix}1\\0\end{pmatrix}}_{\text{Pump term}}, \tag{4.4}$$

where $\Delta_{\mathrm{L/p}} = \omega_{\mathrm{L/p}} - \omega_{\mathrm{R}}$ are the frequencies of the left-cavity resonance and the probe field, respectively, relative to the resonance frequency ω_{R} of the right-cavity mode[1]. The

[1]The reasoning for choosing ω_{R} as the reference is discussed in Section 4.3

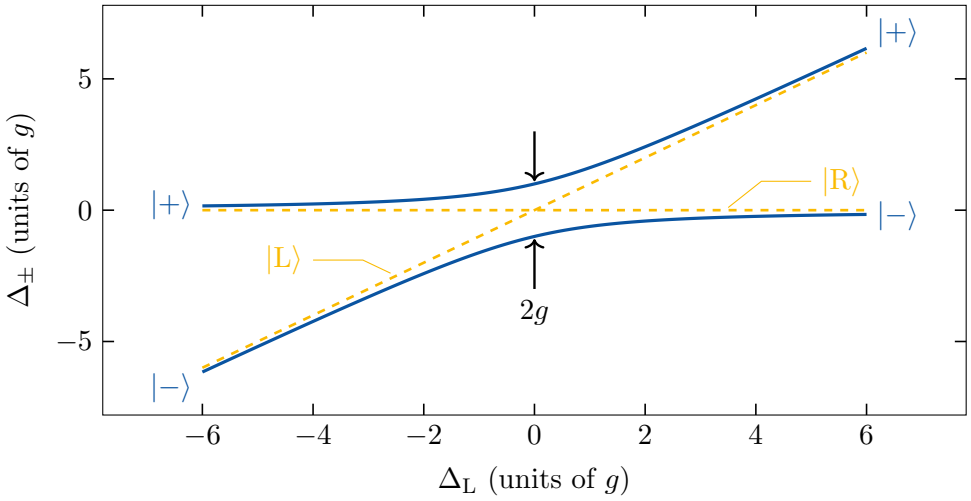


Figure 4.2.: Resonances of two uncoupled (dashed) and coupled (solid) cavities as a function of the cavity–cavity detuning $\Delta_\mathrm{L} = \omega_\mathrm{L} - \omega_\mathrm{R}$. $|\pm\rangle$ denote the symmetric and anti-symmetric superposition of the individual cavity modes $|\mathrm{L/R}\rangle$, and demonstrate mode hybridisation near the centre. The coupling strength, dictating the minimal splitting of the mode branches, is denoted by g.

pump rate $F^\star$ is here set to match the dissipation rate ($\kappa_\mathrm{L}/2$), resulting in normalised field amplitudes [64].

The resonance frequencies of the coupled cavity system correspond to the eigenvalues of the resulting matrix $\tilde{\mathcal{H}}$ representing the Hamiltonian in the semi-classical limit. To find the centre frequencies of the resonances, $\Delta_\pm$, it is sufficient to consider $\kappa_\mathrm{L/R} = 0$, so that the eigenvalues are given by:

$$\Delta_\pm = \frac{\Delta_\mathrm{L}}{2} \pm \sqrt{\left(\frac{\Delta_\mathrm{L}}{2}\right)^2 + g^2}. \tag{4.5}$$

Finally, the corresponding eigenvectors $|\pm\rangle$ of the matrix $\tilde{\mathcal{H}}$ are the symmetric and anti-symmetric superpositions of the individual cavity states $|\mathrm{L}\rangle = (1,0)^T$ and $|\mathrm{R}\rangle = (0,1)^T$:

$$|\pm\rangle = \frac{1}{\sqrt{2}}\left(|\mathrm{L}\rangle \pm |\mathrm{R}\rangle\right), \tag{4.6}$$

as also numerically verified by *transfer-matrix* calculations (see Appendix A.2). These superpositions show that, near $\omega_\mathrm{L} \approx \omega_\mathrm{R}$, the eigenmodes become delocalised across both sub-cavities, forming hybridised normal modes that motivate adopting a photonic–molecule picture [16].

Figure 4.2 visualises this hybridisation by plotting the coupled-cavity resonances $\Delta_\pm$ as functions of the left-cavity detuning Δ_L. The dashed lines indicate the uncoupled resonances, where tuning the left cavity shifts only its own mode while the right-cavity

resonance remains unchanged. In contrast, the solid lines show that coupling lifts the degeneracy near $\omega_\mathrm{L} \approx \omega_\mathrm{R}$, giving rise to the characteristic avoided crossing. At exact degeneracy, $\omega_\mathrm{L} = \omega_\mathrm{R}$, the hybrid modes split symmetrically with a minimal separation of $2g$.

4.1.1. Mode splitting

Generally, the reflectance and transmittance of the coupled cavity system are determined by the coupling to both intracavity modes and can be expressed as:

$$\mathcal{R} = 1 - \left(|\alpha_\mathrm{L}|^2 + |\alpha_\mathrm{R}|^2\right), \tag{4.7a}$$

$$\mathcal{T} = |\alpha_\mathrm{R}|^2, \tag{4.7b}$$

where $|\alpha_\mathrm{L}|^2$ and $|\alpha_\mathrm{R}|^2$ are obtained from Equation (4.4). Since transmission only probes the leakage of the right sub-cavity, one generally finds $\mathcal{R}$ and $\mathcal{T}$ to be unequal even when the system is strongly coupled. The remaining imbalance, quantified by $1 - (\mathcal{R} + \mathcal{T})$, is attributed to non-radiative intracavity losses such as scattering and absorption. In the idealised lossless limit, a strongly detuned second cavity would effectively act as a reflective termination that impedance-matches poorly to the first cavity, forcing the incident field to be almost completely reflected and thereby suppressing excitation of the coupled system except near $\omega_\mathrm{L} \approx \omega_\mathrm{R}$.

In the steady-state limit, where $\partial_t \alpha_{\mathrm{L/R}} = 0$, the reflectance at the degeneracy point ($\omega_\mathrm{L} = \omega_\mathrm{R}$) can be modelled by:

$$\mathcal{R}(\Delta_\mathrm{L} = 0, \Delta_\mathrm{p}) = 1 - \frac{\left(\frac{\kappa_\mathrm{L}}{2}\right)^2 \left(g^2 + \Delta_\mathrm{p}^2 + \left(\frac{\kappa_\mathrm{R}}{2}\right)^2\right)}{g^4 + 2g^2\left(\frac{\kappa_\mathrm{L}\kappa_\mathrm{R}}{4} - \Delta_\mathrm{p}^2\right) + \left(\Delta_\mathrm{p}^2 + \left(\frac{\kappa_\mathrm{L}}{2}\right)^2\right)\left(\Delta_\mathrm{p}^2 + \left(\frac{\kappa_\mathrm{R}}{2}\right)^2\right)}. \tag{4.8}$$

If the coupling strength g, and the dissipation rates κ_L and κ_R, are comparable in magnitude, the resonances become indistinct due to finite linewidths. Conversely, in the strong coupling regime ($4g^2 \gg \kappa_\mathrm{L}\kappa_\mathrm{R}$), the separation between the resonances exceeds their individual linewidths, resulting in the formation of a mode splitting. In this regime, the reflectance can be approximated as:

$$\mathcal{R}(\Delta_\mathrm{p}) \approx 1 - a_g \left(\frac{(\kappa_g/2)^2}{(\Delta_\mathrm{p} - g)^2 + (\kappa_g/2)^2} + \frac{(\kappa_g/2)^2}{(\Delta_\mathrm{p} + g)^2 + (\kappa_g/2)^2}\right), \tag{4.9a}$$

where

$$a_g = 2\left(\frac{\kappa_\mathrm{L}}{\kappa_\mathrm{L} + \kappa_\mathrm{R}}\right)^2 \quad \text{and} \quad \kappa_g = \frac{\kappa_\mathrm{L} + \kappa_\mathrm{R}}{2}. \tag{4.9b, c}$$

Thus, the optical response is characterised by two distinct cavity resonances symmetrically detuned by $\pm g$ from the degeneracy point. However, these resonances exhibit reduced amplitude a_g and modified linewidth κ_g, reflecting dissipative imbalances of the

individual cavities.

The resonance splitting can be qualitatively attributed to the additional phase shift acquired by the circulating field in one cavity upon reflection from the other. The symmetry of this phase shift around resonance conditions enables the constructive interference requirement to be satisfied at two distinct frequencies, giving rise to the observed resonance splitting.

Noteworthy, a similar expression can be derived for the transmittance, which also exhibits two Lorentzian lineshapes with identical centre frequencies and linewidths. However, the transmittance amplitude is further reduced by a factor of two, as transmission arises solely from field leakage through the second cavity, whereas reflectance involves coupling to both cavity modes simultaneously.

4.1.2. Estimation of the coupling strength

The dynamics of the coupled cavity system are fully governed by the resonance frequencies and linewidths of the individual cavities, as well as the coupling strength g. In Chapter 2, we established the relationship between these optical properties and experimentally accessible parameters — such as mirror separation and transmittance, enabling reasonable predictions for these quantities. To effectively design the system parameters, it is necessary to derive a corresponding expression for the coupling strength.

The coupling strength g quantifies the rate at which photons are exchanged between the two cavity modes. Intuitively, a photon circulating within a cavity interacts with the membrane mirror at a rate determined by its round-trip time, given by $c/(2n\cdot d_{\mathrm{L/R}})$, where c/n represents the speed of light in the intracavity medium of refractive index n, and $2d_{\mathrm{L/R}}$ corresponds to the round-trip path length. Each time the photon encounters the membrane, the probability of transmission into the adjacent cavity is proportional to the field transmission amplitude t_{M}. As a result, the coupling strength between two cavities of (approximately) equal length can be approximated as:

$$g \approx \frac{c}{2n \cdot d_{\mathrm{L/R}}} \cdot t_{\mathrm{M}} = \nu_{\mathrm{FSR}} \cdot \sqrt{\mathcal{T}_{\mathrm{M}}}, \tag{4.10}$$

where ν_{FSR} is the free spectral range of the sub-cavity, and $\mathcal{T}_{\mathrm{M}}$ is the transmittance of the membrane. This expression can also be analytically derived in the limit of low membrane transmittance [64].

In the presented experimental setup, the end mirrors are separated by approximately $d \approx 300$ µm, with the membrane positioned at the midpoint, forming two sub-cavities of length $d_{\mathrm{L/R}} \approx 150$ µm. The membrane transmittance, characterised in Section 3.2, is $\mathcal{T}_{\mathrm{M}} \approx 450$ ppm. Using these values, the theoretical coupling strength is calculated as $g_{\mathrm{Theo}} = 2\pi \times 3.38$ GHz.

A comparison of this value with the later measured linewidths of the individual cavities, $\kappa_{\mathrm{L}} = 2\pi \times 212(2)$ MHz and $\kappa_{\mathrm{R}} = 2\pi \times 168(3)$ MHz, as listed in Table 4.1, shows that the system operates within the strong coupling regime, where coherence dominates dissipation ($4g^2 \gg \kappa_{\mathrm{L}}\kappa_{\mathrm{R}}$).

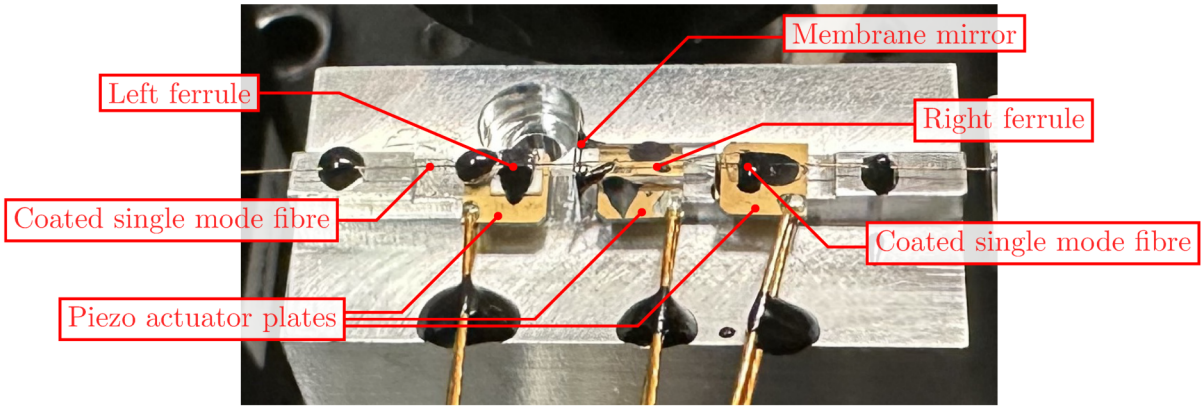


Figure 4.3.: Photograph of the assembled coupled cavity setup. The fibre-based cavity is passively aligned through two different glass ferrules. The right glass ferrule also acts as mount for the membrane mirror. Moreover, the shown arrangement of piezo actuator plates allows for individual addressing of fibre and membrane positions.

4.2. Experimental implementation of the mirror-in-the-middle setup

In Chapter 3, we developed a membrane mirror with an effective diameter of $b_{\mathrm{M}} \approx 96$ µm and a radius of curvature of $R_{\mathrm{M}} \approx 8$ mm. To realise a "mirror-in-the-middle" configuration, this mirror must be integrated into a cavity geometry in which its curvature remains negligible compared to that of the outer cavity mirrors, such that two stable sub-cavities can be formed on either side. We therefore employ a fibre-based cavity system, whose mirrors exhibit radii of curvature more than an order of magnitude smaller than that of the membrane mirror (see Section 2.3), enabling stable plano-convex operation.

The cavity stability criterion (see Section 2.2.3) imposes short cavity lengths for such small radii of curvature. Since the membrane is embedded within a silicon frame of thickness 200 µm [66], the available cavity lengths are comparable to the frame dimensions and leave limited space for integration. To maintain alignment while enabling safe membrane insertion, we pre-align the optical axis of the fibre cavity using two glass ferrules. These ferrules provide passive alignment of the fibre tips and allow the fibres to be retracted, creating sufficient transverse access for inserting the membrane mirror [21]. In addition, the ferrule on the membrane side serves as a mounting interface for attaching the membrane frame close to the cavity axis, thereby improving mechanical stability.

A photograph of the assembled setup is shown in Figure 4.3, where the silicon frame carrying the membrane mirror is mounted directly onto the end facet of the right ferrule.

The full assembly is fixed to an aluminium holder using Stycast epoxy[2], visible as black droplets in the image. The holder is machined with stepped height levels to accommodate the ferrules and piezoelectric actuators while maintaining a common optical axis. Piezoelectric actuators[3] enable axial control of the left fibre position (z_{L}), the right fibre position (z_{R}), and the membrane/ferrule position (z_{M}), allowing fine tuning and modulation of the sub-cavity resonances. The actuators provide a stroke of approximately ± 1.5 µm, enabling tuning across multiple free spectral ranges.

To preserve independent control of all three optical elements, the left fibre is displaced by actuating the left ferrule, whereas on the right side the ferrule (carrying the membrane) and the fibre are actuated independently via separate piezo plates. Since the piezoelectric actuators require high voltages ($U_i = \pm 320$ V), they are electrically insulated from the aluminium mount using alumina spacers.

The construction process required (i) robust passive pre-alignment of the fibre cavity, (ii) controlled membrane insertion without damaging the fibre mirrors, and (iii) fixation of the target sub-cavity lengths while maintaining alignment. The following subsections summarise the key steps and associated practical considerations.

1. Ferrule selection and modification

The ferrules were selected to fulfil two roles: guiding the fibres with passive alignment and providing a mounting surface for the membrane mirror frame. The ferrule used to support the membrane mirror is shown on the right-hand side of Figure 4.4.(a). While its end facet provides sufficient area for bonding the membrane frame, the fibre guide tolerances are tight. For coated fibre mirrors, this becomes relevant because the dielectric film increases the effective fibre diameter near the tip due to the angled vapour deposition process, which coats both the end facet and the sidewalls.

To insert a coated fibre into this ferrule, we therefore removed excess coating from the fibre sidewalls near the tip. This was achieved by gently polishing the fibre on an alumina lapping sheet (grit size 3 µm), drawing it backward while rotating it to ensure uniform abrasion. The fibre was subsequently cleaned in an ultrasonic bath filled with acetone. After several iterations, the fibre could be inserted reliably. Importantly, a cavity test against an opposing fibre mirror showed no measurable degradation of the mirror performance compared to pre-treatment measurements.

To reduce the risk of fibre damage during insertion, a second ferrule with a slightly larger fibre guide was used where membrane support was not required (left-hand side of Figure 4.4.(a)). This ferrule enables straightforward insertion of coated fibres without requiring prior abrasion at the fibre tip. Since its end facet is too small to support the membrane frame, the final design employs both ferrule types: the larger guide ferrule on the side without membrane mounting, and the standard ferrule on the membrane side.

Bonding the membrane frame to the ferrule end facet forms a small enclosed volume between frame and ferrule. Under vacuum conditions, trapped air would slowly escape

[2] LOCTITE® STYCAST 2850FT from Henkel
[3] CSAP02 from Noliac (now CTS Corporation)

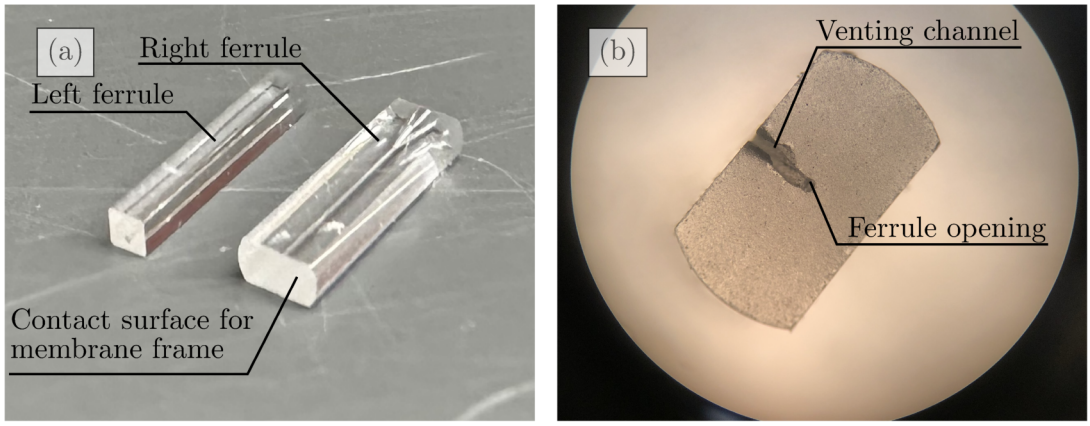


Figure 4.4.: (a) Size comparison of the two glass ferrule types used to guide the left and right fibres. Ferrules with larger end facets provide sufficient area to support the membrane frame, but feature tighter guide tolerances that require additional treatment of coated fibre tips prior to insertion. (b) Microscope image of the end facet of the right ferrule, showing the added venting channel introduced to prevent the formation of a trapped air pocket during vacuum operation. The channel was cut into the glass using a thin diamond wire.

through the fibre guide channel, potentially generating a strong pressure gradient across the membrane at the start of the evacuation process. To mitigate this, a small venting channel was introduced at the ferrule end facet using a thin diamond wire, visible in Figure 4.4.(b). Within the scope of this thesis, the setup was not operated under vacuum, and the venting performance was therefore not tested experimentally.

2. Passive cavity alignment

Small imperfections in fibre mirror fabrication can lead to slight tilts of the mirror surface or lateral offsets between the mirror centre and the fibre core [51–53]. As a consequence, passive alignment alone cannot guarantee optimal cavity mode matching, but it provides a reproducible starting point for fine alignment.

To establish a common optical axis, a single uncoated fibre was threaded through both ferrules. This assembly was then placed onto the respective piezoelectric actuators mounted on the aluminium holder, and the ferrules were oriented such that the fibre remained straight and level. To prevent adhesive or dust from entering the vent opening (see Figure 4.4.(b)), the membrane-side ferrule was mounted with the vent facing downward.

The membrane-side ferrule was fixed first. For this step, the ferrule was held using a tweezer mounted on a five-axis translation stage, where the clamping force was controlled via a fine-thread screw. Stycast epoxy was applied to bond the ferrule to the piezo actuator, and the adhesive was cured for at least 12 hours before proceeding.

After fixation of the first ferrule, the uncoated fibre was removed and replaced by the two coated fibres. Insertion was performed using additional translation stages and USB microscopes for alignment with the ferrule guide. The second ferrule was then aligned

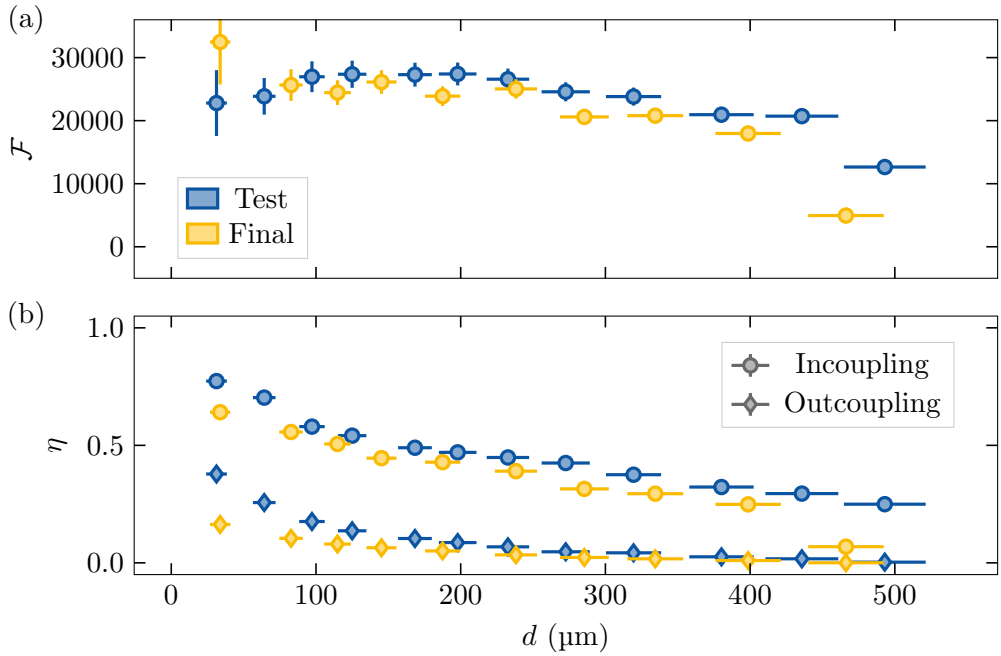


Figure 4.5.: Comparison of (a) Finesse $\mathcal{F}$ and (b) coupling efficiency η of the cavity mode as a function of the cavity length d, before and after passively aligning the fibre–fibre cavity, using two glass ferrules. The corresponding data is presented with blue and orange marker, respectively. Moreover, in (b) the circular and diamond marker correspond to the in- and outcoupling efficiencies, respectively.

by monitoring the cavity formed between the fibre tips and optimising the coupling via small transverse and angular adjustments. Once optimal alignment was achieved, the second ferrule was also bonded to its actuator using Stycast epoxy.

To verify that the ferrule-based mounting preserves cavity performance, we repeated the finesse and coupling-efficiency measurements introduced in Section 2.3. The corresponding results are shown in Figure 4.5. Compared to the original fibre test setup, the ferrule-aligned cavity exhibits comparable finesse and coupling over the target operating range. For larger separations, the finesse decreases at slightly shorter lengths and the coupling efficiency is reduced by a few percentage points. Nevertheless, around the target length $d \approx 300$ µm the ferrule-based alignment reliably reproduces the original performance and allows removal and reinsertion of the fibres without significant degradation.

3. Process of membrane insertion

Following passive alignment, both fibres are fully retracted into their ferrules to protect the fibre mirrors and to create clearance for inserting the membrane frame. This retraction provides a spacing of approximately 2 mm between the ferrules. Since direct optical access to the cavity region becomes obstructed during insertion, the fibre retraction dis-

tances are recorded to enable approximate restoration of their original positions while avoiding collisions with the membrane.

The membrane mirror frame is positioned using the tweezer-mounted five-axis translation stage. Guided by USB microscope images, the membrane is aligned such that its backside (i.e., the side with the etched cavity) is centred with respect to the ferrule axis so that the membrane mirror is located approximately 200 µm from the ferrule opening. Coarse angular alignment is achieved by referencing the ferrule end facet. The membrane frame is then brought into contact with the ferrule, and the fibre is advanced incrementally toward the membrane.

Once a dip in the reflected signal indicates cavity resonance, the alignment is refined by small adjustments of the membrane angle and transverse position. After maximising coupling, the opposing fibre is brought close to the membrane to verify cavity formation on the second side. With confirmed resonances in both sub-cavities, both fibres are retracted and the membrane frame is fixed to the ferrule using small Stycast epoxy drops at the left and right contact points.

4. Fixation of sub-cavity lengths

After the adhesive had fully cured, we removed the tweezer mount and re-evaluated the performance of the two fibre–membrane cavities by measuring their finesse as a function of cavity length. For the left sub-cavity, the cavity length was estimated using USB microscope images. For the right sub-cavity, where direct visual access is obstructed by the membrane frame, the length was determined via the free spectral range extracted from two simultaneously resonant laser fields. Linewidth measurements were calibrated using EOM-generated frequency markers. A detailed description of the measurement procedures is provided in Sections 2.3 and 3.3.

The measured finesse values are shown in Figure 4.6 for two configurations: with the opposing fibre fully retracted (round markers) and with the opposing fibre positioned to form a cavity of length 150 µm on the other side of the membrane (diamond markers). In the first case, the results are consistent with the independent characterisation of both membrane orientations presented in Section 3.3. In particular, the cavity facing the silicon nitride side exhibits substantially higher losses, whereas the opposite cavity shows negligible losses but experiences a faster decline in finesse at larger lengths, consistent with clipping at the fibre mirror.

When both fibres simultaneously form cavities, the linewidths, and thus the inferred losses, are significantly reduced. We quantify this reduction using Equation (3.6) together with the measured finesses and the independently determined mirror transmittances (Section 2.3) and absorption losses of the silicon nitride layer (Section 3.3). Within experimental uncertainties, both sub-cavities show a reduction in effective losses of ~ 335 ppm.

This behaviour indicates that transmission through the membrane mirror no longer acts as an irreversible loss channel when a second cavity is present. Instead, light transmitted through the membrane is coherently recovered by circulating in the opposite sub-cavity, and dissipation occurs predominantly through transmission at the outer fibre

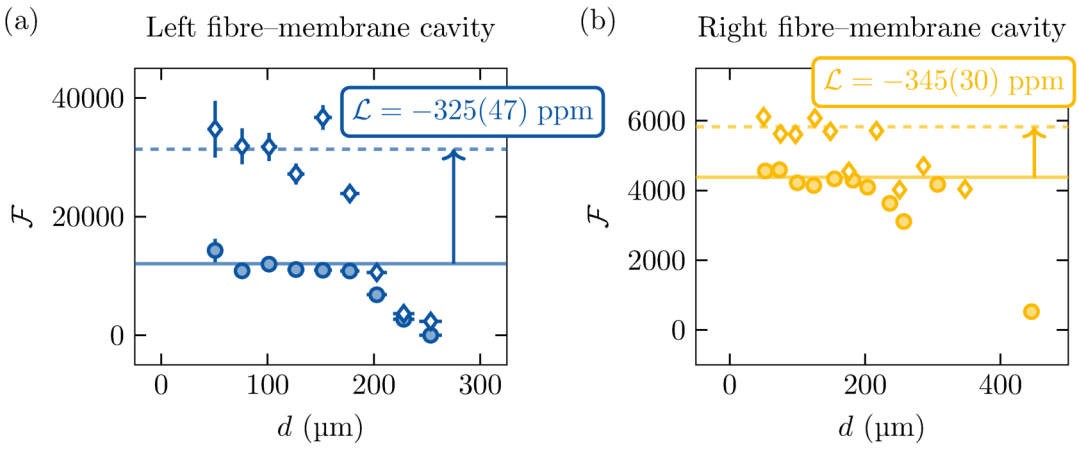


Figure 4.6.: Finesse of the (a) left and (b) right sub-cavities as a function of the respective cavity length. Measurements were first performed with the opposing fibre fully retracted (solid markers), and subsequently repeated with the opposing fibre in place, i.e. in the coupled-cavity configuration. In the latter case, a pronounced increase in finesse is observed, indicating a reduction of effective radiative losses due to coupling through the membrane mirror.

mirrors. We therefore interpret this reduction as evidence for coherent coupling between the two sub-cavities, where membrane transmittance mediates photon exchange rather than introducing effective loss.

Guided by these measurements, both fibres were fixed at the target sub-cavity lengths of $d \approx 150$ µm, where alignment and cavity performance were optimal (see Figure 4.6). The fibres were bonded sequentially, with the respective fibre supported by the tweezer-mounted translation stage to prevent axial displacement during glue application. After curing, the procedure was repeated for the opposite fibre.

Finally, additional epoxy was applied behind the fibre contact points as a strain-relief measure. This reduces the risk of damage to fibres or ferrules due to accidental mechanical contact, e.g. during fibre coupling. In Figure 4.3, these strain-relief droplets correspond to the two outermost Stycast epoxy deposits.

4.2.1. Final sub-cavity performance

With the fibres securely glued to the aluminium mount, the cavity parameters are effectively fixed, with the exception of limited length adjustments achievable via the piezoelectric actuators. These parameters are summarised in Table 4.1, where they are also compared to values measured after the initial operation. Notably, following the initial operation of the setup, the finesse of the left cavity experienced a sudden and significant decline. This decrease coincided with a failure of the laboratory air conditioning system, which resulted in a substantial rise in ambient temperature. Initially, this correlation suggested thermal deformation as a possible cause. However, returning the environment to its original temperature did not restore cavity performance, suggesting that thermal effects alone were unlikely to be the primary factor.

Table 4.1.: Summary of the coupled-cavity parameters. After assembly, the finesse of the left sub-cavity decreased noticeably overnight and remained stable thereafter. For the initial characterisation of the left sub-cavity, the cavity length was estimated visually, whereas subsequent measurements were based on a direct determination of the free spectral range using two frequency-detuned lasers. The same "free-spectral-range method" was used for both measurements of the right sub-cavity.

		After finishing setup	After initial operation
$\nu_{\mathrm{FSR,L}}$	(THz)	0.99(6)	1.02(0)
d_{L}	(µm)	151(8)	147.41(3)
$\delta\nu_{\mathrm{L}}$	(MHz)	28(1)	212(2)
$\mathcal{F}_{\mathrm{L}}$	(10^3)	37(5)	4.8(1)
$\nu_{\mathrm{FSR,R}}$	(THz)	1.01(0)	1.01(0)
d_{R}	(µm)	148.65(3)	148.67(3)
$\delta\nu_{\mathrm{R}}$	(MHz)	177(3)	168(3)
$\mathcal{F}_{\mathrm{R}}$	(10^3)	5.7(1)	6.0(1)

An alternative explanation is that the high circulating power within the pumped cavity may have damaged the membrane mirror, leading to additional optical losses. This hypothesis is supported by the observation that higher-order modes exhibited relatively lower losses. While coupling to higher-order modes is typically suppressed due to the fibre mode's preference for the fundamental Gaussian mode, we observed that the coupling efficiency of the Gaussian mode decreased to a level comparable to that of some higher-order modes. As illustrated in Figure 2.5, the intensity maximum of the Gaussian mode coincides with the intensity minima of many higher-order modes. This suggests that when coupling to higher-order modes, less light interacts with a potentially damaged region of the membrane, reducing its impact on optical losses.

While these explanations remain speculative, it is important to note that no further degradation in cavity performance was observed throughout the remainder of this study. Therefore, the parameters listed in the second column of Table 4.1 remain valid for the setup.

4.3. Spectroscopy and characterisation near degeneracy point

The measurements in Section 4.2 provide initial evidence for coupling between the two sub-cavities: when resonators are formed on both sides of the membrane, the effective cavity lifetime increases. This behaviour is consistent with the fact that an off-resonantly driven cavity can exhibit an effective reflectivity exceeding that of its individual mirrors, thereby reducing the net transmission through the membrane.

Since the membrane transmittance exceeds the transmittance of the outer fibre mirrors, the system operates in a regime where both sub-cavities influence each other in the vicinity of their resonances. As discussed in Section 4.1, this mutual interaction

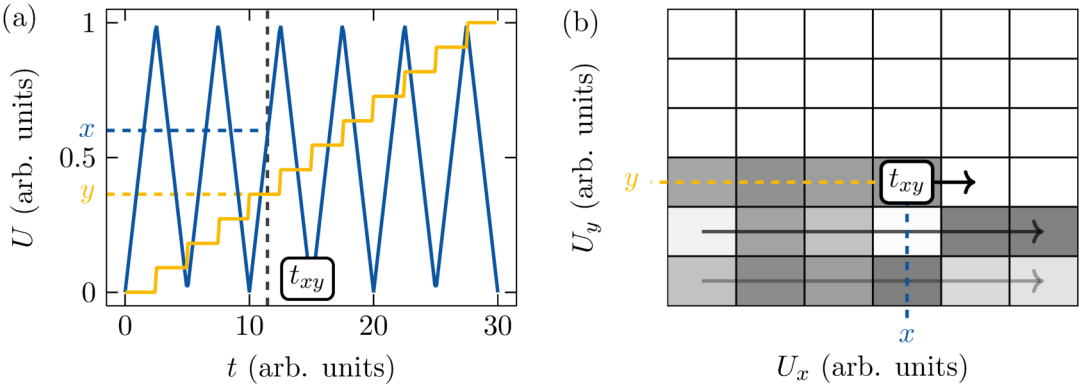


Figure 4.7.: (a) Sketch of the waveforms for scanning ω_L, ω_R or ω_p to resolve the two-dimensional spectrum of the system. (b) Rasterised remapping of the scan voltages. The grey scale of the pixel represent the photo diode signal at the specific moment in time. The established waveform results in a row-by-row population of the image.

gives rise to an avoided crossing, i.e. a lifting of the degeneracy when the uncoupled cavity resonances would otherwise intersect (see Figure 4.2). Consequently, the optical response of the system is not determined by the probed cavity alone, but depends on the detuning of the adjacent cavity as well. Resolving the coupled-cavity spectrum therefore requires systematic control of two independent detuning, yielding a two-dimensional spectral response. In the following, we outline the measurement scheme and the calibration procedures used to obtain high-resolution avoided-crossing spectra.

4.3.1. Scheme of measurement

The coupled-cavity response is governed by three frequencies: the resonance frequencies of the left and right cavities, ω_L and ω_R, and the probe laser frequency ω_p. To map the avoided crossing, two of these quantities must be scanned while the third provides a reference. We implement this by driving two modulations simultaneously, as sketched in Figure 4.7.(a): a fast periodic scan (blue) and a slow ramp (orange). The resulting time trace can be interpreted as a two-dimensional data set, where each timestamp t_{xy} corresponds to a coordinate pair (x, y) and the spectrum is populated row-by-row (Figure 4.7.(b)).

Experimentally, this is realised using continuous waveforms rather than discrete voltage steps. By choosing widely separated timescales, the slow ramp remains approximately constant during one period of the fast scan, such that all samples acquired within one fast cycle can be assigned the same y-coordinate (given by the average value of the slow ramp during this interval). A scan frequency of 600 Hz for the fast waveform and 1 Hz for the slow waveform was found to provide sufficient resolution while keeping the acquisition time short compared to cavity drifts and laser-frequency fluctuations. Due to hysteresis in the piezoelectric actuators, only rising scan segments are evaluated to ensure good alignment of rows in the reconstructed spectrum.

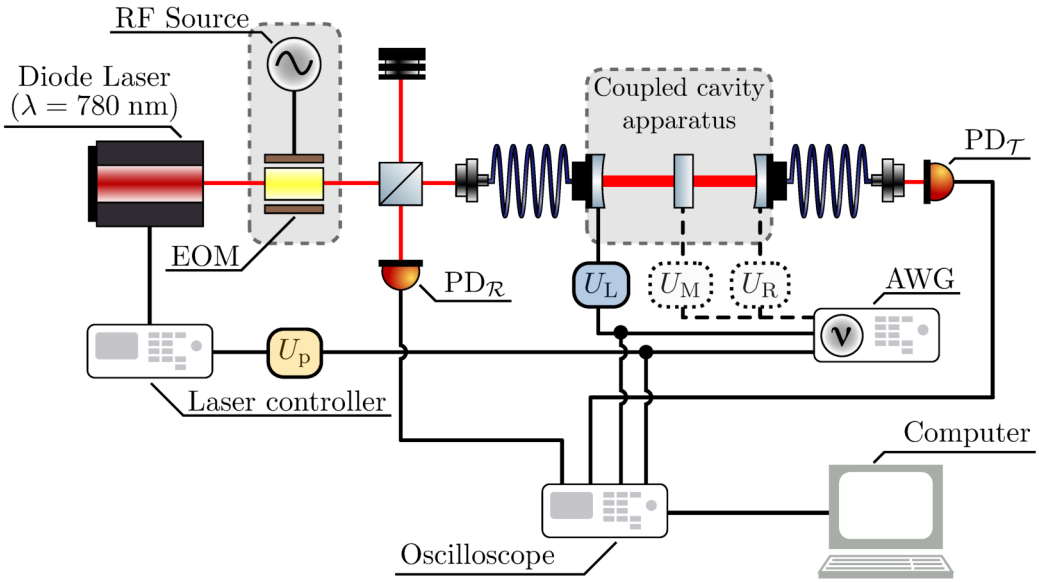


Figure 4.8.: Schematic overview of the optical setup used to acquire two-dimensional reflection and transmission spectra of the coupled-cavity system. The spectrum is mapped by scanning a two-parameter space through simultaneous modulation of two frequency scales, chosen from ω_L, ω_R, and the probe laser frequency ω_p, using an arbitrary waveform generator (AWG). The photodiode signals are recorded and reconstructed into a two-dimensional spectrum in post-processing. Optionally, an electro-optic modulator (EOM) is driven to generate sidebands that provide frequency markers for accurate calibration of the scan axes. This figure includes elements generated with Ref. [54].

Experimental implementation

A schematic of the optical setup is shown in Figure 4.8. A diode laser at $\lambda \approx 780$ nm probes the coupled-cavity system in reflection and transmission. Frequency markers are generated using an electro-optic modulator (EOM), and scan signals are applied using an arbitrary waveform generator (AWG). The reflected and transmitted signals, together with the scan voltages, are recorded on an oscilloscope and processed on a computer to reconstruct the two-dimensional spectrum. The basic optical layout and linewidth extraction procedure follow the approach introduced in Section 2.3.

In addition to scanning cavity length, we also employ frequency modulation of the probe laser. The laser frequency is tuned via a piezo actuator acting on the grating inside the laser cavity, supplemented by feed-forward current modulation to suppress mode hops [98]. This feed-forward control is provided by a commercial laser controller[4], which interfaces between the AWG and the laser.

In the measurements presented here, we modulate the left-cavity resonance ω_L and the probe laser frequency ω_p. This choice is motivated by the presence of strong friction between the right fibre and the ferrule housing the membrane mirror, which introduces

[4] DLC pro: Digital Control from TOPTICA Photonics AG

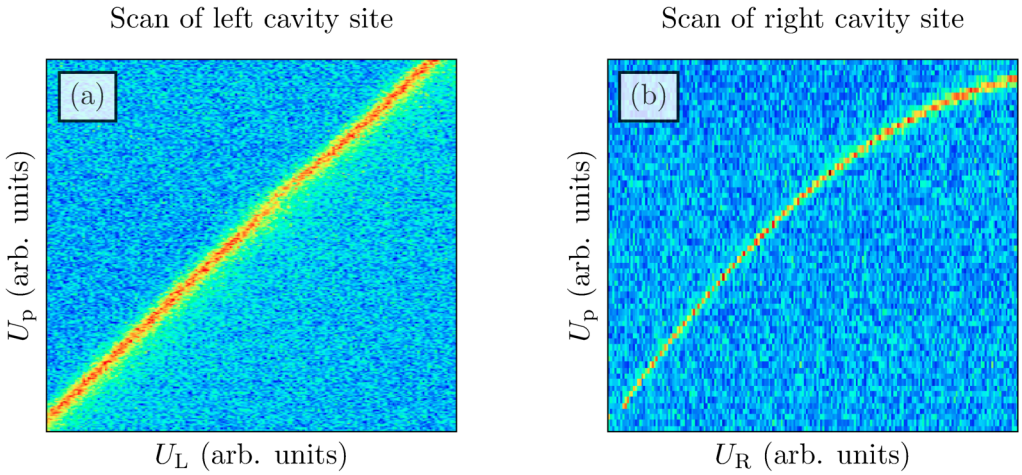


Figure 4.9.: Two-dimensional reflection spectrum of the (a) left and (b) right sub-cavities as a function of the piezo actuator voltage that controls the respective cavity length. The voltage that controls the laser frequency is shown vertically in both figures.

a pronounced nonlinearity between the applied scan voltage and the resulting resonance detuning. This effect is illustrated in Figure 4.9, where separate two-dimensional scans of the reflected field of the left and right cavities are compared.

Furthermore, we apply the fast modulation to ω_L while scanning ω_p slowly. This ordering reduces the impact of slow power variations during laser tuning, which arise from residual beam displacement during piezo tuning, current modulation, and etalon effects in the coupling path. From the periodicity of the observed intensity modulation, we infer an effective etalon length of approximately 32 cm, consistent with the length of the coupling fibre and suggesting reflections between its end facets. By scanning the laser slowly, the optical power is approximately constant over one cycle of the cavity-length modulation. During post-processing, the DC component is therefore removed row-by-row to suppress residual power offsets.

However, a residual linear background remains in the reflected signal. If the incoupling efficiency were independent of the cavity detuning, this background could be removed by fitting a straight line to the measured resonance amplitudes and subtracting it from the data. Near the avoided crossing, this assumption breaks down: the hybridisation of the two cavity modes reduces the maximal coupling efficiency (see Equation (4.9)), so the resonance amplitudes no longer trace the underlying linear background alone. To obtain a reliable estimate of the background gradient, we therefore restrict the linear fit to the edges of the scanned frequency range, where the cavities are far detuned from each other. In this limit, the cavity–cavity coupling is negligible and the maximal incoupling can again be treated as constant. In our measurements, the scan range can exceed $2\pi \times 15$ GHz, while the maximum expected coupling strength is only $g = 2\pi \times 3.38$ GHz, ensuring that the scan boundaries lie well outside the coupling region and can be used to determine and remove the linear gradient.

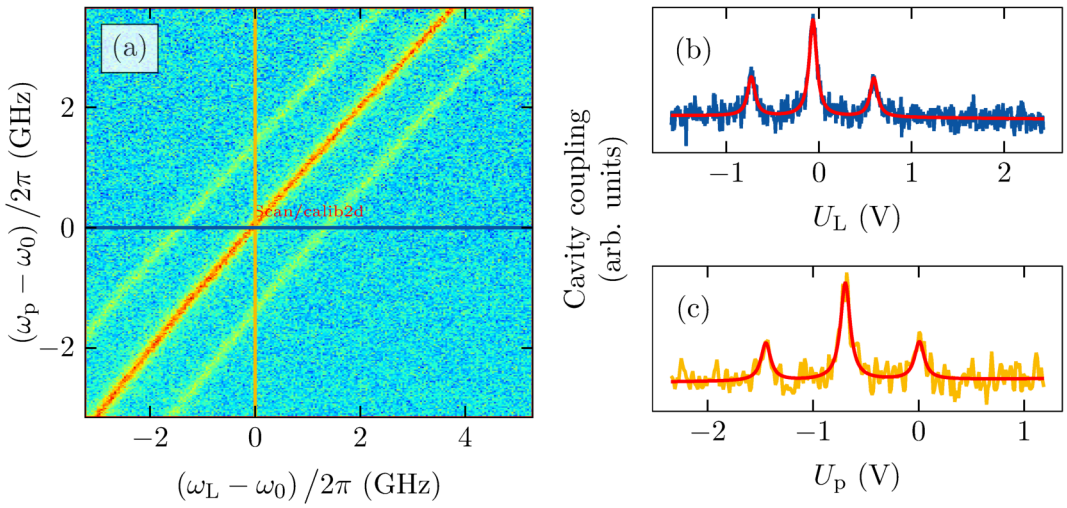


Figure 4.10.: (a) Two-dimensional measurement of the (inverted) reflection spectrum of the left cavity, with sidebands modulated on the incident probe beam, resulting in three parallel resonances, detuned by ±1.4 GHz. Subfigures (b) and (c) correspond to the horizontal and vertical cuts, indicated in (a) and show three Lorentzian lineshapes that are used to translate voltage scales to frequencies.

Calibration of scan axis

To convert scan voltages into frequency units, we use the EOM-generated sidebands as calibration markers. In the two-dimensional spectrum, these sidebands produce three parallel resonance traces (carrier and $\pm\Omega$ sidebands), separated by the known modulation frequency (e.g., here ±1.4 GHz). An example calibration measurement is shown in Figure 4.10.(a), recorded with the right cavity far detuned. The corresponding horizontal and vertical cuts, indicated in the figure, are fitted with three Lorentzian profiles to determine the voltage-to-frequency conversion along both axes.

Applying this calibration and tracking the carrier resonance reveals that the two scan axes are matched, with a slope of 1.001(2) GHz/GHz, consistent with the expected 1:1 detuning relation. The extracted linewidths, $\kappa_x = 2\pi \times 206(11)$ MHz and $\kappa_y = 2\pi \times 210(17)$ MHz, also agree within uncertainty and are consistent with prior measurements of the left-cavity linewidth κ_L (after its degradation).

At this stage, the calibrated axes remain referenced to an arbitrary origin defined by the chosen cuts through the spectrum. To locate the right-cavity resonance frequency ω_R in the following avoided-crossing measurements, we fit the extracted resonance centres to the coupled-mode dispersion relation given in Equation (4.5).

Finally, since the piezoelectric response depends on both scan amplitude and modulation frequency, the calibration procedure was repeated regularly to ensure consistent frequency scaling.

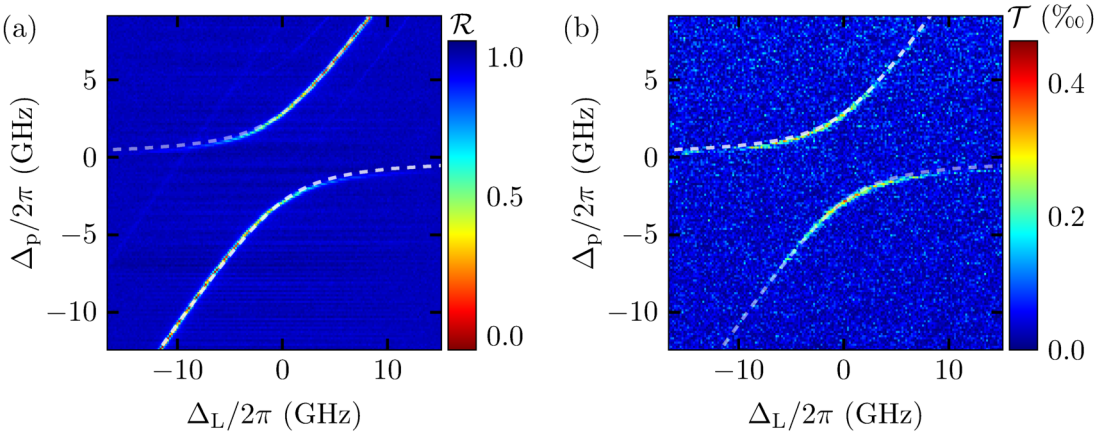


Figure 4.11.: Two-dimensional (a) reflection and (b) transmission spectra of the coupled-cavity system, near the degeneracy point, where $\omega_L \approx \omega_R$. The system is probed from the left side, resolving an avoided crossing. The dashed lines represent a fit to the trace of the observed minima in reflectance, mirroring the dispersion relation of the system.

4.3.2. Measurement of the avoided crossing

To locate the degeneracy point at which both cavities are simultaneously near resonance, we probe the left cavity by scanning its length and recording the reflected signal. As the right cavity is tuned towards resonance, the left-cavity response changes markedly in centre frequency, linewidth, and amplitude. In addition, once the right cavity approaches resonance, light can couple through the system and a measurable transmission signal can be observed.

Figure 4.11 shows the optical response of the coupled-cavity system in the vicinity of this degeneracy point. Subfigure (a) displays the reflected intensity measured at the left cavity, while Subfigure (b) shows the corresponding transmission signal detected on the right-hand side. The right-cavity resonance frequency ω_R is obtained by fitting the measured resonance centres to the coupled-mode dispersion relation given in Equation (4.5). The resulting fit is indicated by dashed lines in both subfigures.

The measured avoided crossing illustrates the evolution of the eigenstates $|\pm\rangle$ as they asymptotically approach the individual cavity states $|L\rangle$ and $|R\rangle$ when moving away from the avoided crossing. As indicated in Figure 4.2, each coupled state initially corresponds to a single cavity mode but undergoes mixing at the avoided crossing, evolving with opposite assignments thereafter. This behaviour is evident in the spectrum, where coupling to the incident light requires at least partial overlap with the left-cavity mode. Consequently, as the horizontal line at $\Delta_p = 0$ is approached, the right cavity becomes fully resonant with the light field. However, the left cavity suppresses the field, resulting in negligible coupling. Conversely, the transmission spectrum exhibits a symmetric evolution, as transmission requires coupling to both modes, leading to compensation of asymmetries. Ultimately, the transmission signal vanishes when no light can couple into either cavity, whereas the reflectance remains. This behaviour indicates that

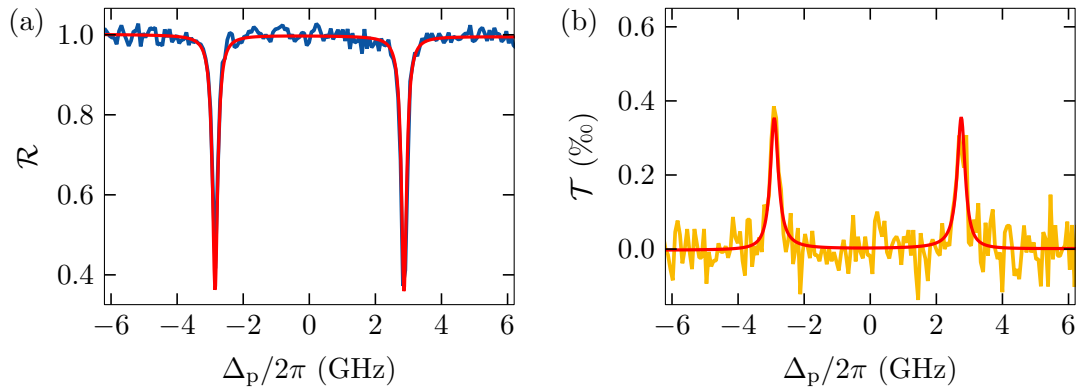


Figure 4.12.: (a) Normalised reflectance and (b) normalised transmittance of the coupled-cavity system as a function of the probe-laser detuning Δ_p, extracted from the vertical cuts in Figure 4.11. For this measurement the two sub-cavities are tuned into mutual resonance ($\Delta_L = 0$), such that the spectrum resolves the minimal splitting of the hybridised normal modes.

the incoupling through the input mirror is still balanced by other loss channels: once the transmission channel through the right cavity is effectively closed, this balance is provided predominantly by non-radiative intracavity losses.

The cavity–cavity coupling strength g is obtained from the minimum splitting between the two branches of the avoided crossing. In the measured spectrum, the smallest separation occurs at $\Delta_L = 0$. Vertical cuts through the reflection and transmission spectra at this detuning are shown in Figure 4.12.

Both traces exhibit two well-resolved resonances with comparable linewidths. Fitting the reflection spectrum with a model consisting of two Lorentzian line shapes yields a mode splitting of $2g = 2\pi \times 5.718(4)$ GHz, a linewidth of $\kappa_g = 2\pi \times 189(6)$ MHz, and a peak amplitude of $a_g = 0.64(3)$. Using the independently measured cavity linewidths $\kappa_L = 2\pi \times 212(2)$ MHz and $\kappa_R = 2\pi \times 168(3)$ MHz, Equation (4.9) predicts $\kappa_g \approx 2\pi \times 190$ MHz and $a_g \approx 0.62$, in excellent agreement with the extracted values.

From the measured splitting we obtain a coupling strength of $g = 2\pi \times 2.859(2)$ GHz, which is slightly lower than the theoretical prediction $g_{\mathrm{Theo}} = 2\pi \times 3.38$ GHz. A plausible explanation is imperfect mode overlap between the two sub-cavities, e.g. due to a residual misalignment of their optical axes, which reduces the probability that light transmitted through the membrane couples efficiently into the opposing cavity mode and thereby lowers the effective coupling.

Consistent with this interpretation, the transmitted signal is substantially weaker than expected from the reflected spectrum, indicating imperfect coupling to the output fibre mode. As a consequence, the transmission amplitude cannot be interpreted as a direct measure of the right-cavity population $|\alpha_R|^2$. In addition, the linewidth extracted from the transmission trace is significantly larger, $\kappa_g = 2\pi \times 360(40)$ MHz, which we attribute to the reduced signal-to-noise ratio and the limited bandwidth of the photodetector at the required gain settings. Importantly, the extracted mode splitting remains consistent with the value obtained from reflection, supporting the robustness of the coupling estimate.

4.4. Summary and future improvements

In this chapter, we demonstrated the realisation of a coupled-cavity device based on a ferrule-aligned fibre microcavity. One ferrule additionally serves as a mechanical support for the membrane mirror, resulting in a compact and mechanically robust assembly. By combining piezoelectric tuning with a two-dimensional scanning scheme, implemented via two simultaneous frequency sweeps and subsequent post-processing, we obtained high-resolution spectra of the coupled system. These measurements reveal a clear avoided crossing, providing direct evidence of strong coupling between the two sub-cavities.

The observed avoided crossing can be interpreted within the photonic-molecule picture, where the coupled cavity resonances form hybridised normal modes that are formally analogous to those encountered in light–matter systems. In this sense, the device provides a purely optical analogue of the characteristic signatures of linear light–matter interaction, a viewpoint that will be developed further in the next chapter [15, 16].

From the minimum mode splitting, we extract a coupled-mode linewidth and amplitude of $\kappa_g = 2\pi \times 189(6)$ MHz and $a_g = 0.64(3)$, respectively. Both values agree well with theoretical expectations, validating the underlying model and confirming the reliability of the experimental analysis. The coupling strength is measured as $g = 2\pi \times 2.859(2)$ GHz, approximately 20 % below the predicted value, which is consistent with residual mode mismatch between the two sub-cavity modes.

While the present implementation already operates within the expected parameter regime, several improvements could further enhance its performance and flexibility. A modified ferrule design with a slightly wider fibre guide and a larger end facet would reduce mechanical friction and enable wider-range scans of the right fibre or the membrane position. In addition, replacing the single-mode output fibre with a multimode fibre could increase the detected transmission signal. Although single-mode fibres provide optimal spatial mode matching to the cavity [9], the same strict mode selectivity also suppresses transmission whenever the output coupling is imperfect. A multimode fibre would relax this constraint, improving collection efficiency of the transmitted signal, but at the expense of direct coupling efficiency to the right-cavity mode.

More broadly, the use of multimode fibres could also facilitate coupling to higher-order transverse cavity modes via controlled mode distortions. This trade-off is particularly relevant in the context of the following chapter, where higher-order modes are explored as a tool for tuning the cavity–cavity coupling strength. In the present configuration, the use of single-mode fibres prevents a direct exploration of these effects and necessitates indirect measurement strategies.

Regardless, the realised setup already constitutes a stable and versatile platform for investigating coupled dissipative cavity systems.

5. Tuneable coupling strength by higher-order mode mixing

Contents

5.1. Mode overlap of higher-order transversal cavity modes . . . 73
5.2. A purely optical analogue to the Purcell effect 75
5.2.1. Dissipative interaction in the intermediate and strong coupling regime 78
5.3. Formation of the minimal mode splitting 81
5.4. Range of accessible coupling strengths 83
5.5. Simulation of coupled-resonator-induced-transparencies . . . 85
5.6. Summary 88

In Chapter 4 we integrated a membrane mirror into an existing fibre-cavity resonator, thereby forming two coupled sub-cavities. Spectroscopy revealed a clear avoided crossing with a minimal mode splitting near zero detuning, i.e. when both sub-cavities are resonant to each other. The observed hybridisation reflects coherent photon exchange between the two sub-cavities and is reminiscent of characteristic mode-mixing phenomena known from linear light–matter interaction. Motivated by this correspondence, we can interprete the coupled system as a photonic molecule, i.e. a pair of optical resonances whose hybridised normal modes provide a purely optical analogue of key aspects of light–matter physics [15, 16]. In this chapter we develop this viewpoint further and demonstrate how the photonic molecule can be tuned to study different dissipative regimes.

To access these different regimes, we must be able to tune the system parameters, namely the cavity–cavity coupling and the dissipation rates. It is precisely this control over the parameters that sets a platform apart as a useful laboratory for open quantum systems: by sweeping the balance between coherent exchange and loss, we can observe how states, spectra, and linewidths are shaped by an open environment [1–3, 5, 99].

However, directly varying the intrinsic dissipation rates is not possible here, as they are set by the end-mirror transmittances and by scattering and absorption losses, which are not tuneable in a controlled manner. Fortunately, the relevant dynamics are governed not by dissipation alone but by the ratio of coupling strength to dissipation. In quantum optics this ratio, the cooperativity, quantifies the competition between coherent exchange and decay. Consequently, different dissipative regimes can be reached by tuning the coupling strength instead. The minimum coupling is set by the membrane's

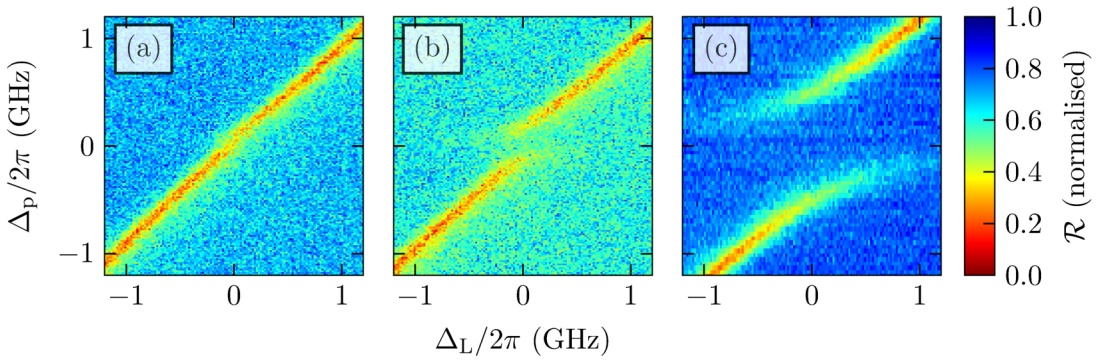


Figure 5.1.: Series of two-dimensional reflection spectra measured on the left cavity for different higher-order cavity modes resonant with the right cavity, showing different degrees of an avoided crossing. The spatial overlap between the two cavity modes increases progressively, demonstrating (a) weak, (b) intermediate, and (c) strong coupling, respectively.

finite transmittance, implying a lower bound on the achievable coupling. To maintain sufficient mode coupling and optical access, this transmittance is typically chosen comparable to the end-mirror transmittance (here, $\mathcal{T}_M \approx 450$ ppm and $\mathcal{T}_F \approx 120$ ppm, respectively), which would make the weak-coupling regime inaccessible.

However, by restricting the spatial overlap of the two cavity modes, the coherent exchange between the sub-cavities can be reduced substantially. In this chapter we show that a slight misalignment of the optical axes (consistent with the observations in Section 4.3.2) can lift the degeneracy of transverse cavity modes and enabling a tool for controlled reduction of the intercavity mode overlap. Figure 5.1 presents a series of reflection spectra, obtained using the two-dimensional measurement scheme introduced in Section 4.3.1, in which the fundamental TEM_{00} mode of the left cavity is tuned into resonance with higher-order TEM modes of the right cavity. From left to right, the spatial overlap increases progressively, and the avoided crossing evolves accordingly. In particular, the near-vanishing splitting in Subfigure (a) demonstrates that the effective coupling can be reduced by almost two orders of magnitude, thereby accessing the weak-coupling regime.

We use this mechanism to characterise the system in the weak-coupling limit, revealing clear analogies to the Purcell effect [100], a hallmark of weak light–matter interaction, further reinforcing the photonic-molecule picture. Because the overlap control is discretely set by addressing specific transverse-mode pairs, we additionally survey and quantify a wide range of coupling strengths, observing the gradual formation of an avoided crossing and the emergence of a minimal mode splitting, as already hinted at in Figure 5.1.

Restricting the mode overlap affects the two sub-cavities symmetrically. Therefore, while it provides access to both weak- and strong-coupling regimes, it does not offer independent control over the dissipation rates of the two resonators. Since higher-order

transverse modes exhibit only marginally increased losses, this tuning knob is insufficient for engineering a pronounced dissipation imbalance. To access regimes dominated by asymmetric decay, we instead modulate the two cavity lengths at different rates during the scan. We show that this effective dissipation imbalance enables coupled-resonator-induced-transparency (CRIT)-like interference phenomena in the optical response, where distinct decay pathways can destructively interfere and give rise to narrow transparency features reminiscent of electromagnetically induced transparency (EIT) [14, 101].

5.1. Mode overlap of higher-order transversal cavity modes

The cavity–cavity coupling strength g quantifies the rate of coherent photon exchange between the optical modes of the two sub-cavities. Accordingly, g is not determined solely by the transmittance of the shared membrane mirror $\mathcal{T}_\mathrm{M}$ and the cavity lengths (via $\nu_\mathrm{FSR} \propto d^{-1}$, see Equation (4.10)), but also by the spatial overlap of the transverse field distributions on the membrane. In particular, photons transmitted through the membrane contribute to coherent coupling only if they efficiently match the optical mode of the adjacent cavity. Otherwise, they do not populate the opposite cavity mode and are effectively dissipated.

This dependence can be expressed by writing the effective coupling strength as:

$$\begin{aligned} g &= \left(\nu_\mathrm{FSR}\sqrt{\mathcal{T}_\mathrm{M}}\right)\cdot\eta \\ &= \left(\nu_\mathrm{FSR}\sqrt{\mathcal{T}_\mathrm{M}}\right)\cdot\left(\frac{|\int \mathrm{d}S\boldsymbol{\mathcal{E}}_\mathrm{L}(r,\phi,0)\boldsymbol{\mathcal{E}}_\mathrm{R}^*(r,\phi,0)|^2}{\int \mathrm{d}S|\boldsymbol{\mathcal{E}}_\mathrm{L}(r,\phi,0)|^2\cdot\int \mathrm{d}S|\boldsymbol{\mathcal{E}}_\mathrm{R}(r,\phi,0)|^2}\right)^{\frac{1}{2}}. \end{aligned} \tag{5.1}$$

where $\eta \in [0,1]$ denotes the normalised transverse mode overlap of the left and right intracavity fields evaluated across the membrane surface. Here, $\boldsymbol{\mathcal{E}}_\mathrm{L}(r,\phi,0)$ and $\boldsymbol{\mathcal{E}}_\mathrm{R}(r,\phi,0)$ represent the transverse field distributions of the two sub-cavities at their focal plane. For the approximately plano–concave sub-cavity geometries realised in this work, the membrane is located close to the beam waist, such that focal plane and membrane mirror position coincide. Here, the fields are not restricted to the fundamental Gaussian TEM_{00} mode but may correspond to higher-order Laguerre–Gaussian TEM_{lp} modes introduced in Section 2.2.3.

For perfectly aligned and symmetric cavities, the transverse cavity modes form an orthonormal basis [43], however. Consequently, the overlap vanishes for non-identical mode indices, $\{l,p\} \neq \{l',p'\}$, such that $\eta = 0$ unless the transverse modes on both sides of the membrane are identical. This orthogonality is illustrated in Figure 5.2.(a), where the bar heights show the overlap between a given TEM_{lp} mode and TEM_{00} mode.

In the present setup, the experimentally observed coupling between the two TEM_{00} modes is lower than predicted (see Section 4.3.2), consistent with a residual misalignment between the optical axes of the two sub-cavities. A simple model captures this effect by introducing a lateral displacement δr of one mode relative to the other. Matching the observed reduction in coupling by approximately 20 % corresponds to $\delta r \approx$

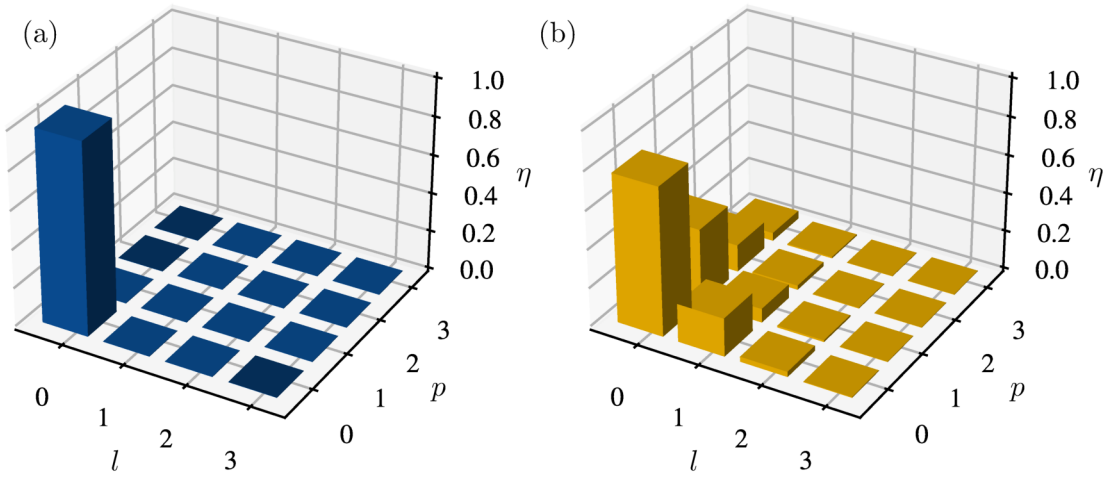


Figure 5.2.: Mode overlap η of the TEM_{00} mode with higher-order Laguerre–Gaussian modes, characterised by the indices l and p, for (a) the mode centres perfectly aligned and (b) the mode centres shifted by $\delta r \approx 0.7\, w_0$ so that the reduced measured coupling strength is replicated.

5.33 µm $\approx 0.7\, w_0$. Such a displacement lifts the strict orthogonality of the transverse eigenmodes and yields non-zero overlap between modes with different $\{l, p\}$, as shown in Figure 5.2.(b).

This mechanism provides a practical route to tune the effective coupling strength. By selectively addressing higher-order transverse modes in one sub-cavity, the overlap η, and thus g, can be reduced well below the nominal scale set by $(\nu_{\mathrm{FSR}}\sqrt{\mathcal{T}_{\mathrm{M}}})$. In the fibre-based implementation, however, this tunability comes with an experimental constraint, as the single-mode fibres act as spatial mode filters. Therefore, direct excitation of non-Gaussian modes is strongly suppressed, thereby reducing coupling to the external probe field [9]. By tuning only the second cavity into higher-order modes, this suppression can be largely confined to the transmission side of the system, while leaving the reflected signal of the probed cavity essentially unaffected.

Quantitative prediction of the coupling strength requires reliable identification of the transverse modes involved. Direct imaging of the intracavity intensity distributions is not feasible in the present architecture, as coupling into the fibre waveguides distorts the spatial mode profiles. Instead, transverse modes can be identified spectroscopically, since higher-order modes acquire an additional round-trip phase via their Gouy phase contribution, given by:

$$\psi_{lp}(z_{1/2}) = (l + 2p + 1)\psi(z_{1/2}), \tag{5.2}$$

where $\psi(z_{1/2})$ denotes the Gouy phase of the underlying Gaussian mode at the fibre mirror positions, given by Equation (2.11d). In principle, this additional phase shifts the resonance frequencies of higher-order transverse modes relative to TEM_{00} and thereby

encodes the mode indices in the additional detuning.

In practice, unambiguous assignment of (l, p) from resonance frequencies is challenging in the present cavity geometry. Due to the relatively large Rayleigh length, the Gouy-phase-induced spacing between adjacent transverse mode families is comparable to the free spectral range ($\psi(z_{1/2}) \approx 200$ GHz versus $\nu_{\mathrm{FSR}} \approx 1$ THz). As a result, transverse and longitudinal modes are not cleanly separated and can spectrally overlap, such that multiple combinations of longitudinal and transverse indices may appear at similar resonance frequencies. This mixing makes unique identification of l and p based solely on the resonance frequencies practically infeasible.

5.2. Two weakly coupled cavities: A purely optical analogue to the Purcell effect

By coupling the fundamental TEM_{00} mode in the actively probed left cavity to a higher-order mode in the right cavity, we access a regime of weak cavity–cavity coupling, which manifests as a local modification of the left cavity's reflectance around $\Delta_{\mathrm{L}} \approx 0$, as exemplified in Figure 5.1.(a).

In the context of light–matter interaction, this regime is described by the Purcell effect [100], where the emission rate of an emitter is modified by the presence of a resonant cavity. In our system, the left cavity plays the role of the "emitter", while the right cavity modifies its effective emission rate — or loss — by controlling the fraction of light transmitted through the shared mirror. Equivalently, the coupled-cavity system can be viewed as a single cavity in which one end mirror is replaced by a secondary, tuneable cavity whose detuning-dependent reflection, transmission, and phase response dress the effective end mirror.

In this picture, the reflectance of the first cavity as a function of its detuning Δ_{L} can be written as a Lorentzian (analogous Equation (2.7)) of the form:

$$\mathcal{R}(\Delta_{\mathrm{L}}) = 1 - a(\Delta_{\mathrm{p}}) \frac{(\gamma(\Delta_{\mathrm{p}})/2)^2}{(\Delta_{\mathrm{L}} - \delta(\Delta_{\mathrm{p}}))^2 + (\gamma(\Delta_{\mathrm{p}})/2)^2}, \tag{5.3}$$

where $\Delta_{\mathrm{L}} = \omega_{\mathrm{L}} - \omega_{\mathrm{R}}$ is the detuning between the two cavity resonances. The dependence on the probe frequency ω_{p} is absorbed into $\delta(\Delta_{\mathrm{p}})$, so that the reflectance is expressed as a function of Δ_{L}, with $\Delta_{\mathrm{p}} = \omega_{\mathrm{p}} - \omega_{\mathrm{R}}$ acting as a tuneable parameter. The parameters that define the lineshape — the full width at half maximum γ, the centre frequency δ, and the amplitude a — all depend on Δ_{p}, the detuning between the incident light field and the resonance of the second cavity.

Rewriting the generalised reflectance (see Equation (4.8), with $\Delta_{\mathrm{L}} \neq 0$) in this form,

the parameters γ, δ, and a can be expressed as:

$$\gamma(\Delta_\mathrm{p}) = \kappa_\mathrm{L}\left(1 + \frac{4g^2}{\kappa_\mathrm{L}\kappa_\mathrm{R}}\left|\tilde{\alpha}_\mathrm{R}(\Delta_\mathrm{p})\right|^2\right), \tag{5.4a}$$

$$\delta(\Delta_\mathrm{p}) = \Delta_\mathrm{p}\cdot\left(1 - \frac{4g^2}{\kappa_\mathrm{R}^2}\left|\tilde{\alpha}_\mathrm{R}(\Delta_\mathrm{p})\right|^2\right), \tag{5.4b}$$

$$a(\Delta_\mathrm{p}) = \frac{\kappa_\mathrm{L}}{\gamma(\Delta_\mathrm{p})}, \tag{5.4c}$$

where $|\tilde{\alpha}_\mathrm{R}|^2$ denotes the normalised lineshape of the right cavity in the absence of coupling ($g = 0$). For the linear response discussed here this lineshape is Lorentzian; analysing the nonlinear response in Chapter 6 suggests that the same formalism remains applicable even when the right-cavity resonance acquires more complex lineshapes.

Experimentally, this effective single-cavity description can be tested by measuring the reflectance $\mathcal{R}$ as a function of Δ_L while the detuning Δ_p is varied. To do so, we modulate only the resonance frequency ω_L of the left cavity and systematically step the probe laser frequency ω_p, thereby changing Δ_p. For each value of Δ_p, we extract the fitted Lorentzian parameters γ, δ, and a. In the two-dimensional reflection spectra established earlier (e.g. Figure 5.1), this corresponds to horizontal cuts through the data.

This analysis is illustrated in Figure 5.3. Subfigure (a) shows the full two-dimensional reflection spectrum of the TEM_{00} mode in the left cavity, weakly coupled to a higher-order mode in the right cavity. As the resonances overlap, clear variations in linewidth and amplitude appear, even though the interaction is too weak to generate a resolvable avoided crossing. Subfigures (b) – (d) display the extracted Lorentzian parameters, centre frequency δ, linewidth γ, and amplitude a, as functions of Δ_p, together with fits based on Equations (5.4a) to (5.4c) (shown in red). The small deviations visible in subfigure (b) are attributed to nonlinear effects in the right cavity, which will be discussed in Chapter 6.

In particular, subfigure (c) shows that the linewidth of the left cavity increases as the probe frequency approaches resonance with the right cavity ($\Delta_\mathrm{p} \approx 0$), directly analogous to the enhanced emission rate of an emitter in a resonant cavity described by the Purcell effect. Not only is the behaviour qualitatively similar to light–matter interaction, but it is described by the same equations. In the traditional Purcell scenario, the enhancement of the emission rate is quantified by the Purcell factor F [102], which is equivalent to the cooperativity and compares the coupling strength to the product of the decay rates of emitter and cavity. Here, Equation (5.4a) yields the corresponding relative linewidth enhancement:

$$F_P = \frac{4g^2}{\kappa_\mathrm{L}\kappa_\mathrm{R}}. \tag{5.5}$$

Fitting the data in Figure 5.3.(c) gives a Purcell factor of $F_P = 0.34(6)$. From the fitted lineshape and background level, we extract $\kappa_\mathrm{R} = 2\pi \times 218(23)$ MHz and $\kappa_\mathrm{L} =$

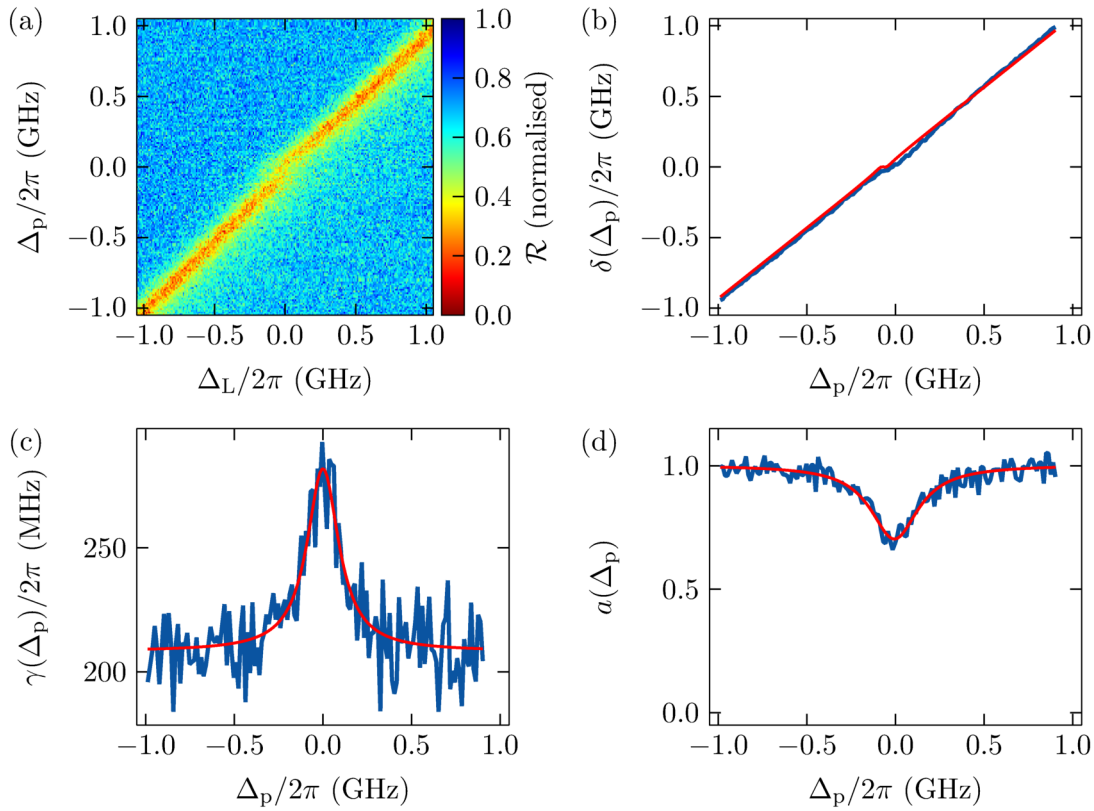


Figure 5.3.: (a) Two-dimensional reflection spectrum of the left cavity, with a mode with only small spatial overlap being resonant with the right cavity, resulting in weak coupling. (b), (c), (d) represent centre frequency δ, linewidth γ, and amplitude a retrieved from a Lorentzian fit on the horizontal profile of (a) as functions of the laser detuning Δ_p. These parameter hide the unperturbed lineshape of the right cavity, best visible in Subfigure (c).

$2\pi \times 208(1)$ MHz, and consequently a coupling strength $g = 2\pi \times 63(8)$ MHz.

This analytic treatment thus enables a complete characterisation of the coupling between the two cavities even in the absence of a clearly resolved avoided crossing or direct spectroscopy of the secondary (right) cavity resonance, which is hindered by its weak coupling to the output waveguide. In this way, the alternative description of the coupled cavity setup reproduces effects known from light–matter interaction in a purely optical environment, motivating the photonic-molecule interpretation and its relevance as an experimentally accessible analogue [15, 16].

5.2.1. Dissipative interaction in the intermediate and strong coupling regime

In light–matter interactions, the Purcell effect is typically associated with the weak-coupling regime, where the interaction between an emitter and a cavity is insufficient to form hybridised modes. One practical reason is that the emitter resonance usually cannot be tuned arbitrarily, so one focuses on regimes where dispersive effects are small and the resonance shift remains negligible, i.e. $\delta(\Delta_\mathrm{p}) \approx \Delta_\mathrm{p}$ [103].

In contrast, the coupled-cavity analogue presented here allows independent tuning of the "emitter" resonance ω_L, so that dispersive and dissipative contributions to the response can be disentangled. Also, the mathematical model introduced above was derived without assuming weak coupling, implying that it remains valid also in the strong-coupling regime. This is demonstrated in Figure 5.4, where the extracted parameters are analysed for a configuration in which the TEM_{00} mode in the left cavity is strongly coupled to a mode in the right cavity, clearly resolving an avoided crossing.

However, stronger coupling leads to a suppression of the resonance amplitude near the degeneracy point, as predicted by Equation (5.4c) and confirmed by Figure 5.4.(d). This becomes problematic when the amplitude suppression causes unresolved inflection points in the Lorentzian lineshape of the right-cavity resonance. In such cases, the amplitude and width of the lineshape cannot be unambiguously determined, which prevents a reliable extraction of the Purcell factor F_P.

To address this, the extracted linewidths $\gamma(\Delta_\mathrm{p})$ can be normalised with respect to the coupling strength g, according to:

$$\frac{\gamma(\Delta_\mathrm{p})}{g^2} = \frac{\kappa_\mathrm{L}}{g^2} + \frac{4}{\kappa_\mathrm{R}}|\tilde{\alpha}_\mathrm{R}|^2 = \frac{\kappa_\mathrm{L}}{g^2} + \frac{\kappa_\mathrm{R}}{\Delta_\mathrm{p}^2 + (\kappa_\mathrm{R}/2)^2}. \tag{5.6}$$

This normalisation inherently links amplitude and width and thus resolves the ambiguity — at the expense of requiring prior knowledge of g. The coupling strength can be extracted independently by exploiting its signature in the resonance shift $\delta(\Delta_\mathrm{p})$, which becomes pronounced in the strong-coupling regime and behaves as:

$$\delta(\Delta_\mathrm{p}) \underset{\Delta_\mathrm{p}, g \gg \kappa_\mathrm{R}}{\approx} \Delta_\mathrm{p} - \frac{g^2}{\Delta_\mathrm{p}}. \tag{5.7}$$

Figure 5.4.(b) shows a fit[1] to the measured resonance shift, which enables a robust extraction of the coupling strength. Combined with the normalised fit in Figure 5.4.(c), this procedure allows us to determine g, κ_L, and κ_R under strong-coupling conditions. For the configuration shown in Figure 5.4, we obtain $g = 2\pi \times 678(48)$ MHz, $\kappa_\mathrm{L} = 2\pi \times 210(15)$ MHz, and $\kappa_\mathrm{R} = 2\pi \times 300(10)$ MHz.

The good agreement between data and fits strongly suggests that this dissipative-coupling picture, in which the secondary cavity modulates the effective dissipation of

[1] The fit shown is based on Equation (5.4a). However, κ_R cannot be reliably extracted from this fit alone, as it mainly affects the slope near the resonance centre, where no data are resolved. In the regions where data are available, the simplified and full models converge.

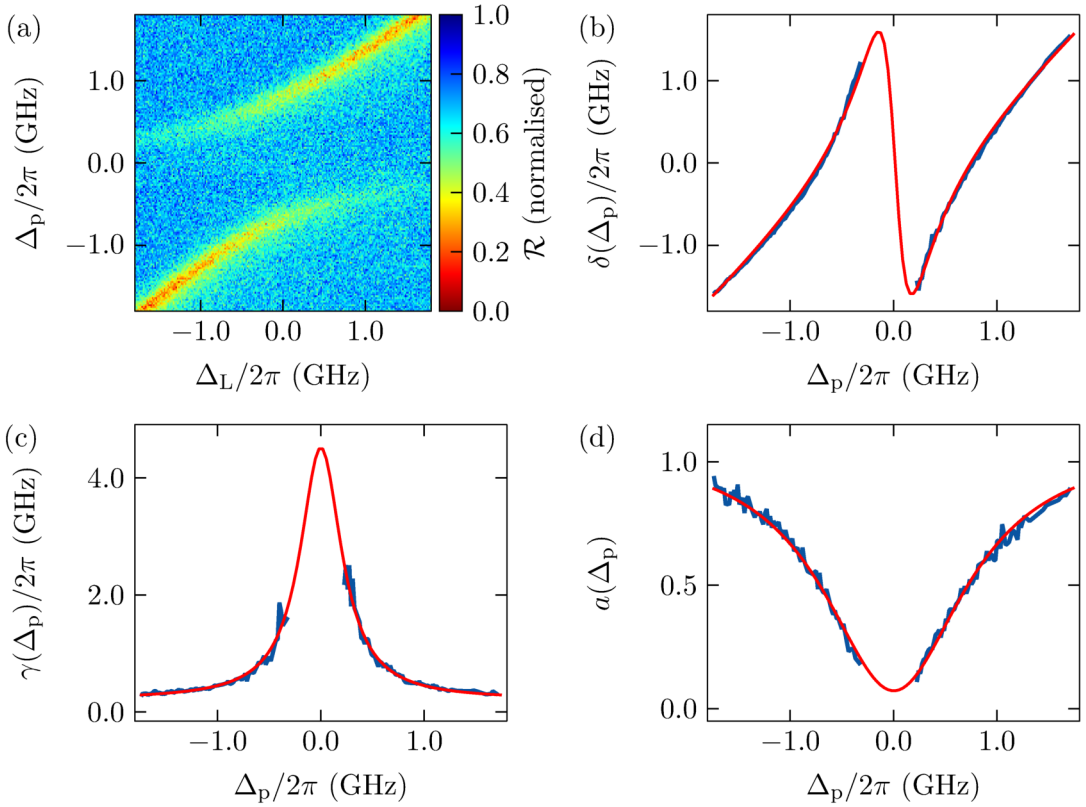


Figure 5.4.: (a) Reflected signal of the left cavity as a function of cavity detuning Δ_L and laser detuning Δ_p, while the TEM_{00} in the left cavity strongly couples to a higher-order cavity mode in the right cavity, as evidenced by the mode splitting. (b), (c), (d) represent centre frequency δ, linewidth γ, and amplitude a retrieved from a Lorentzian fit on the horizontal profile as function of the laser detuning Δ_p. Low optical contrast in the centre region prevents extraction of the parameters. By combining fits to data presented in (b) and (c) the gap can be interpolated.

the primary cavity, remains valid even in the strong-coupling regime. Consequently, this approach provides a robust and versatile method for characterising the coupling strength between the two cavities across a broad range of regimes. This is particularly useful given that we cannot perform direct spectroscopy of the right-cavity resonance: higher-order cavity modes are typically not observable with single-mode fibres [9], as the fibre acts as an efficient spatial mode filter, suppressing coupling to non-Gaussian modes. To overcome this limitation, we employ the horizontal-cut analysis introduced in this section, which allows us to probe the dissipative coupling between the two cavities and to extract the linewidth κ_R of the right-cavity resonance.

Although, in principle, the cavity linewidth is independent of the transverse mode profile, implying constant values for κ_L and κ_R, this assumption strictly holds only for cavities with infinitely large mirrors. In practice, as the transverse mode order increases, the spatial extent of the optical field grows, leading to additional clipping losses at the

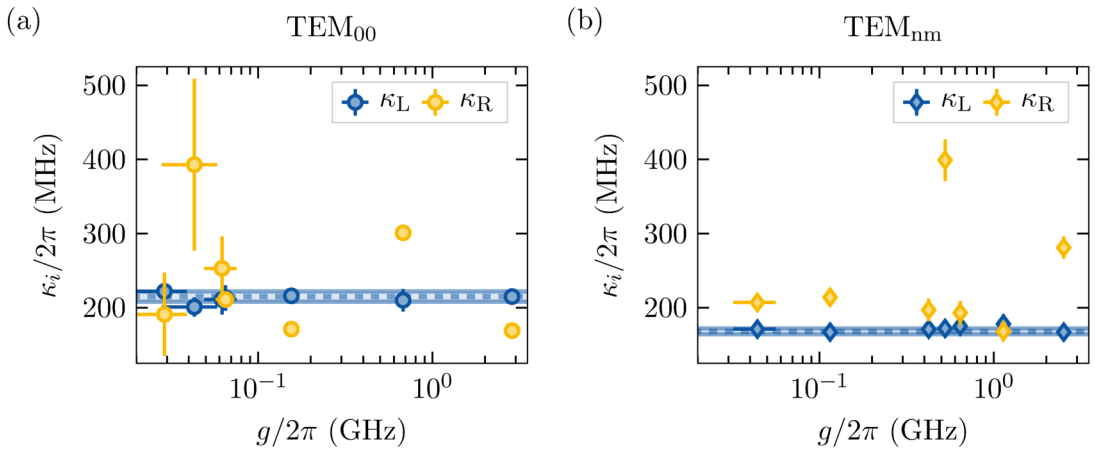


Figure 5.5.: Cavity linewidths κ_L and κ_R as functions of the coupling strength g, for (a) the fundamental TEM_{00} mode and (b) an unknown mode of higher-order in the left cavity. While the respective dissipation rates of the left cavity are consistent across all observed couplings, varying dissipation rates in the second cavity are related to the coupling to different transversal modes, exhibiting various degree on losses.

finite edges of the fibre mirrors, as previously discussed in Section 2.3.2, or on the boundary of the membrane mirror. In the preceding subsections, we examined two distinct transverse modes in the right cavity, each exhibiting different spatial overlaps with the left-cavity mode and indicating different linewidths. It therefore becomes essential to determine κ_R (and κ_L) individually for each observed mode to accurately characterise the system.

Figure 5.5 summarises all extracted linewidths κ_L and κ_R, ordered by their respective coupling strength g, as obtained using the methods established in this section. During the analysis, two distinct modes in the left cavity were found to couple to the fibre waveguide, resulting in two independent data sets, presented in the left and right subfigures, respectively.

These two modes differ in optical contrast (10 % vs. 6 %) and linewidth ($\kappa_L = 2\pi \times 215(8)$ MHz vs. $\kappa_L = 2\pi \times 168(4)$ MHz), suggesting that they correspond to different transverse cavity modes. Based on the mode-selective nature of single-mode fibre coupling, the mode exhibiting stronger coupling is identified as the fundamental TEM_{00} mode. The increased losses observed for the TEM_{00} mode are likely caused by surface damage on the membrane, which induces additional scattering at the beam centre, as discussed in Section 4.2. The second mode exhibits reduced coupling and lower loss, consistent with a higher-order transverse mode with minimum intensity in the centre ($l > 0$, see Figure 2.5), and is therefore referred to as TEM_{nm}[2].

[2] A plausible candidate is the lowest-order mode with suppressed on-axis intensity, TEM_{10}. This assignment is supported by the overlap trends shown in Figure 5.2 and by the approximately Gaussian guided fibre mode, which favours weak coupling to higher-order modes. However, because the transverse profile cannot be imaged directly in the fibre-based geometry, this identification remains tentative.

In both subfigures, the linewidth κ_{L} of the left cavity remains approximately constant. In contrast, the right cavity exhibits a manifold of coupled modes with varying linewidths κ_{R}, without an immediately obvious pattern. Nevertheless, some notable similarities emerge: both left-cavity modes appear to couple to a subset of right-cavity modes with comparable linewidths, suggesting that the same right-cavity modes are involved. For instance, in both cases we can identify one mode with $\kappa_{\mathrm{R}} \approx 2\pi \times 400$ MHz and another with $\kappa_{\mathrm{R}} \approx 2\pi \times 300$ MHz.

In summary, we demonstrate that dissipative coupling between the two cavities can be used to perform indirect spectroscopy of the right sub-cavity, allowing us to resolve its resonance lineshape despite the inaccessibility imposed by single-mode fibre coupling. This analysis is crucial for interpreting the extracted coupling strengths, since the coherent rate g only acquires clear physical meaning when considered together with the associated dissipative rates. Moreover, it provides an essential reference for understanding how g and the two linewidths jointly shape the measured spectra.

5.3. Transition from weak to strong coupling: Formation of the minimal mode splitting

The lifted degeneracy between transverse cavity modes in the two sub-cavities can not only be used to access the weak-coupling regime, but also to demonstrate the formation of the (minimal) mode splitting in discrete steps, as coherence overcomes dissipation. To this end, we evaluate the cavity reflectance $\mathcal{R}$ at the degeneracy point ($\Delta_{\mathrm{L}} = 0$) as a function of the detuning between the probe laser and the right-cavity resonance, Δ_{p}, for different transverse cavity modes being resonant with the right cavity, exhibiting reduced cavity–cavity coupling, caused by the reduced mode overlap with the left cavity mode.

In the strong-coupling regime, this degeneracy point is marked by the minimal spectral separation between the two branches of the avoided crossing. In the weak-coupling regime, however, such a splitting is not developed, making this point inaccessible via conventional avoided-crossing analysis. However, as established in Section 5.2, the lineshape of the right-cavity resonance can be inferred indirectly from its dissipative influence on the left-cavity resonance. In this approach, the degeneracy point is identified via $\gamma(\Delta_{\mathrm{p}})$ and $\delta(\Delta_{\mathrm{p}})$ (see Equations (5.4a) to (5.4b)), allowing its determination even when the splitting cannot be resolved.

Figure 5.6 presents the reflectance at the degeneracy point for all observed couplings, corresponding to vertical slices of the two-dimensional reflectance spectra at identified $\Delta_{\mathrm{L}} = 0$. The two transverse cavity modes supported by the left cavity, designated TEM_{00} and $\mathrm{TEM}_{\mathrm{nm}}$, are shown in the left and right columns, respectively. The data are ordered by increasing coupling strength g, extracted by fitting the measured reflectance to the analytical model derived in Equation (4.8).

In the weak-coupling regime, however, the parameters κ_{L}, κ_{R} and g all contribute to both the amplitude and the linewidth of the resonance, making independent determination of each parameter ambiguous. Attempts to fix κ_{L} based on prior measurements

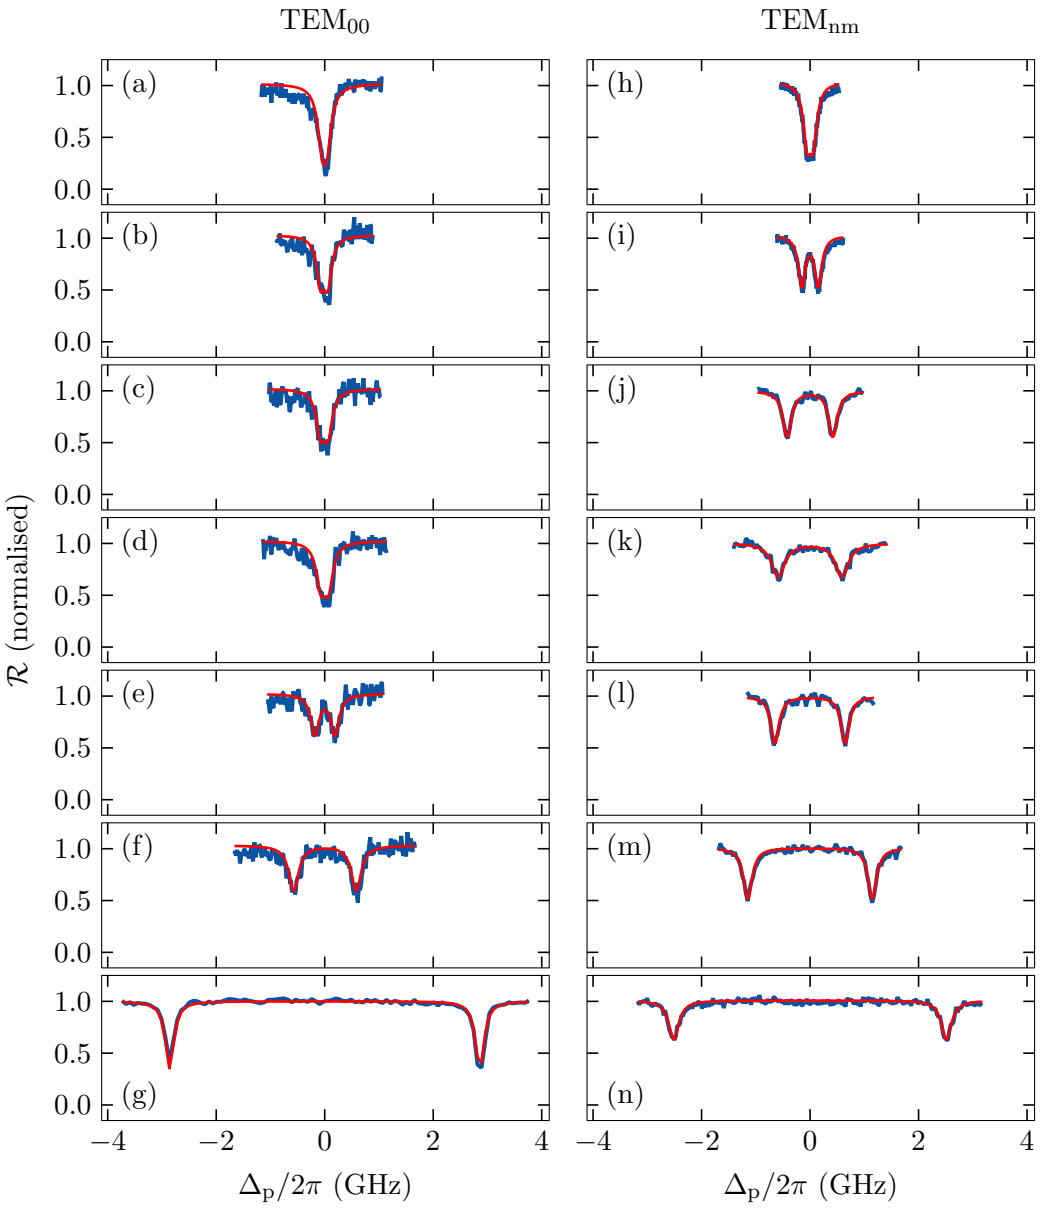


Figure 5.6.: Reflection spectra of the coupled cavities, when the left cavity is resonant with higher-order transverse cavity modes in the right cavity, $\Delta_L = 0$, as functions of Δ_p. As the mode overlap between the resonant modes decreases, the coupling strength g, extracted from the fit (red) also decreases, enabling discrete control over this parameter. The spectra are sorted in ascending order from top to bottom, revealing the formation of two split modes, indicating mode hybridisation. Measurements involving the TEM_{00} mode or the higher-order transversal mode in the left cavity are shown in the left and right column, respectively.

led to unstable fits that were highly sensitive to the initial estimates of the remaining parameters. To address this, we introduce an amplitude parameter a_g as an additional free parameter and combine κ_L and κ_R into an effective linewidth κ_g, thereby reducing the parameter space and stabilising the fit. The simplified analytical model (shown in red) then becomes:

$$\mathcal{R}(\Delta_\mathrm{L}=0,\Delta_\mathrm{p}) = 1 - a_g \frac{\left(\frac{\kappa_g}{2}\right)^2 \left(g^2 + \Delta_\mathrm{p}^2 + \left(\frac{\kappa_g}{2}\right)^2\right)}{g^4 + 2g^2\left(\left(\frac{\kappa_g}{2}\right)^2 - \Delta_\mathrm{p}^2\right) + \left(\Delta_\mathrm{p}^2 + \left(\frac{\kappa_g}{2}\right)^2\right)^2}. \tag{5.8}$$

The initial fit instability arises in part from the sensitivity of the coupled-mode amplitude to the underlying parameters, which is further amplified by residual nonlinearities in the signal introduced during the laser scan. As discussed in Section 4.3.1, scanning the laser frequency creates a power gradient that, despite post-processing correction, leaves small variations in optical power. Introducing a free amplitude parameter effectively compensates for these effects and enables reliable extraction of the coupling strength across the full range of interactions.

With increasing coupling strength g, Figure 5.6 illustrates the progressive formation of the minimal mode splitting: In subfigures (a) and (h), only a single Lorentzian resonance is observed with a broadend linewidth, consistent with the Purcell effect. As the coupling strength g increases, two distinct symmetric resonances emerge, corresponding to the formation of hybridised normal modes. In subfigures (f) – (g) and (k) – (n), this mode separation becomes pronounced, clearly indicating entering strong coupling.

Interestingly, the $\mathrm{TEM_{nm}}$ mode in the left cavity couples more uniformly to the transverse modes in the right cavity, resulting in a more gradual formation of the mode splitting. In contrast, most interactions involving the $\mathrm{TEM_{00}}$ mode lie in the weak-coupling regime, where the response is much more sensitive to small variations of the right-cavity resonance frequency ω_R. Nonlinear effects in the right cavity lead to power-dependent shifts and distortions of ω_R, which therefore show up more prominently in the $\mathrm{TEM_{00}}$ data set. As a consequence, measurements involving the $\mathrm{TEM_{00}}$ mode were carried out at lower probe powers to mitigate these nonlinearities, at the cost of a reduced signal-to-noise ratio. The sensitivity of the extracted parameters to probe intensity, particularly at low coupling strengths, is analysed in more detail in Chapter 6.

5.4. Range of accessible coupling strengths

We established two complementary methods to determine the coupling strength g between the cavity modes. The first is based on analysing the dissipative interaction of the probed cavity with the static secondary cavity, corresponding to horizontal cuts through the two-dimensional reflection spectra. The second relies on mode hybridisation that forms an avoided crossing, extracted from vertical cuts through the spectra. In this section, we briefly confirm the consistency between both approaches and assess the range of accessible coupling strengths, exploiting the lifted degeneracy between the transverse

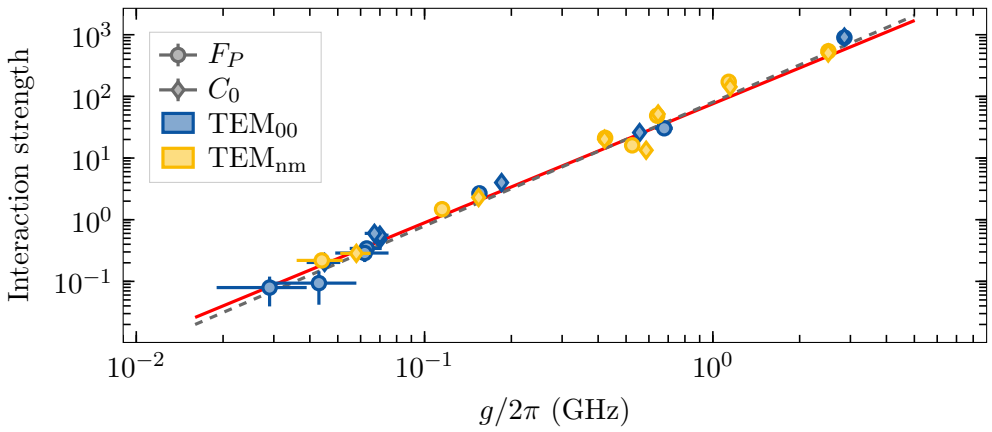


Figure 5.7.: Interaction strength (F_P, C_0) as a functions of the coupling strength g reaching over almost three orders of magnitude, ranging from weak to strong coupling. Results from two different methods, i.e., analysing horizontal and vertical profiles in the two-dimensional reflection spectrum, are displayed with circular and diamond markers, respectively. The colour of the marker indicates the resonant mode in the probed left cavity. All data aligns on a line representing the quadratic relation.

cavity modes. However, the coupling strength g alone provides limited insight unless it is contextualised with respect to the system's dissipation rates κ_L and κ_R, or the effective linewidth κ_g. We therefore also consider the two dimensionless quantities:

$$F_P = \frac{4g^2}{\kappa_\mathrm{L}\kappa_\mathrm{R}} \quad \text{and} \quad C_0 = \left(\frac{2g}{\kappa_g}\right)^2, \tag{5.9}$$

analogous to the Purcell factor F_P and the cooperativity C_0, which characterise the weak- and strong-coupling regimes, respectively.

In Figure 5.7, we plot the dimensionless interaction parameter as a function of the coupling strength g, using data from both analysis methods, indicated by different markers. Measurements involving the fundamental TEM_{00} mode and the higher-order transverse mode $\mathrm{TEM}_{\mathrm{nm}}$ in the left cavity are shown in blue and orange, respectively.

Both expressions in Equation (5.9) predict a quadratic dependence on the coupling strength, indicated by the dashed guideline in the figure. A fit to the data, shown as a solid red line, yields an exponent of $k = 1.91(41)$, in good agreement with the theoretical expectation. It should be noted, however, that the interaction parameters are functionally derived from the fitted coupling strengths within the same model, so an intrinsic correlation is built in. The scatter observed around the ideal trend is attributed to variations in the dissipation rates κ_L, κ_R, or κ_g across different cavity modes. Nonetheless, the data show strong consistency across different mode orders and fitting methods, while clearly illustrating the broad range of experimentally accessible interaction strengths.

In summary, coupling to higher-order transverse modes in the right sub-cavity allows the coupling strength g to be varied over nearly two orders of magnitude. As a result,

the interaction parameter spans the range $F, C_0 \in [10^{-1}, 10^3]$, granting access to both weak- and strong-coupling regimes.

It is worth emphasising that, in principle, the system supports even weaker coupling. The right sub-cavity can accommodate an infinite set of transverse modes, and as illustrated in Figure 5.2, the spatial mode overlap decreases rapidly with increasing mode order. Although higher-order modes may be suppressed by clipping losses at the fibre mirrors, it remains plausible that additional, yet undetected, modes exist within the right cavity. However, below a certain threshold the corresponding coupling strengths would induce changes in amplitude or phase that fall below the detection sensitivity, making such interactions effectively undetectable in practice.

5.5. Dynamic cavity resonance modulations: Simulation of coupled-resonator-induced-transparencies

Electromagnetically induced transparency (EIT) is a prominent example of dissipation engineering [101]. In a three-level atomic system, a control and a probe field drive two coherent excitation pathways whose interference suppresses absorption and opens a narrow transparency window in the spectrum [93]. Phenomenologically, the system response is reshaped by destructive interference between these pathways, giving rise to a dark state that is effectively decoupled from the probe field.

The two coupled cavity modes that form the photonic molecule resemble the coupled atomic transitions. In this picture, mode hybridisation can likewise produce EIT-like spectral signatures in a purely optical environment, commonly referred to as coupled-resonator–induced transparency (CRIT) [14].

The standard EIT/CRIT condition requires (i) sufficiently weak coupling such that the hybridised modes overlap and interfere, and (ii) a strong dissipation imbalance, e.g. $\kappa_R \gg \kappa_L$ [104, 105]. Under these conditions, a narrow transparent feature of width $\Delta\omega = \kappa_L + 4g^2/\kappa_R$ can emerge[3].

In our coupled-cavity platform, condition (i) can be reached by selecting appropriate transverse-mode pairs to access the weak-coupling regime. Condition (ii), in contrast, is not naturally available because the intrinsic dissipation rates of the two sub-cavities remain comparable ($\kappa_R \approx \kappa_L$), even for higher-order transverse modes. To nevertheless access CRIT-like interference physics, we dynamically engineer an effective dissipation imbalance by scanning the two cavity resonances at unequal rates.

Intuitively, the linewidth sets the scale converting detuning into reflectance change; when the detuning evolves in time, the relevant "rate" that shapes the line is the detuning sweep per unit time. Driving ω_L and ω_R with different scan amplitudes therefore mimics unequal linewidths, an effective $\kappa_R \neq \kappa_L$, even if the intrinsic optical losses are similar.

A practical handle for this comes from a mechanical asymmetry identified in Section 4.3.1: friction between the right fibre mirror and the membrane ferrule reduces

[3]The width follows from the standard EIT picture with the substitution $\gamma_{13} \to \kappa_R$, $\gamma_{12} \to \kappa_L$, and $|\Omega_c| \to 2g$, where γ_{13} is the decay rate of the probed transition, γ_{12} is the decay rate of the meta-stable dark state, and $|\Omega_C|$ is the Rabi frequency of the strong control field, see [14, 106].

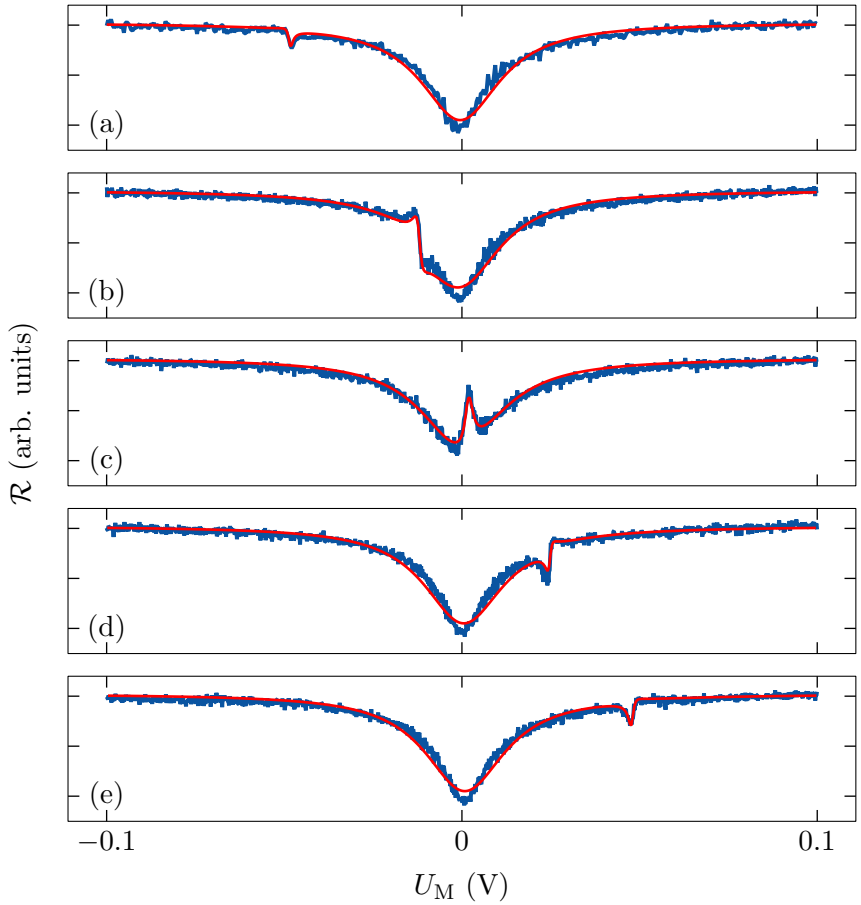


Figure 5.8.: Reflected signal of the right cavity as a function of the piezo actuator voltage controlling the membrane position, U_M, showing the right-cavity resonance. Subfigures (a) to (e) show different detuning between both cavity sites, $\Delta_\mathrm{L} = \omega_\mathrm{L} - \omega_\mathrm{R} \approx 2\pi \times \{-308\ ,\ -78\ ,\ 12\ ,\ 158\ ,\ 306\}$ MHz, revealing a second narrow resonance belonging to the left cavity. Overlap of the resonances leads to destructive interference effects.

the responsiveness of the right cavity. We therefore modulate the membrane position using the designated piezoelectric actuator, driven by the voltage U_M. Translating the membrane simultaneously shortens one sub-cavity while lengthening the other, but the corresponding detuning change with different slopes:

$$\delta_\mathrm{L}(U_\mathrm{M}) = \zeta_\mathrm{L}\, U_\mathrm{M}, \quad \text{and} \quad \delta_\mathrm{R}(U_\mathrm{M}) = -\zeta_\mathrm{R}\, U_\mathrm{M}. \tag{5.10}$$

Calibrating the frequency shift per applied piezo voltage using the measured linewidths of the left and right-cavity resonances, we obtain $|\zeta_\mathrm{R}| \approx 6$ MHz/mV and $|\zeta_\mathrm{L}| \approx 150$ MHz/mV. Thus, $|\zeta_\mathrm{R}/\zeta_\mathrm{L}| \approx 0.04 \ll 1$, which provides the required scan-rate asymmetry and enables synchronous scans with strongly unequal amplitudes.

Figure 5.8 presents the measured and fitted reflectance of the right cavity versus U_M

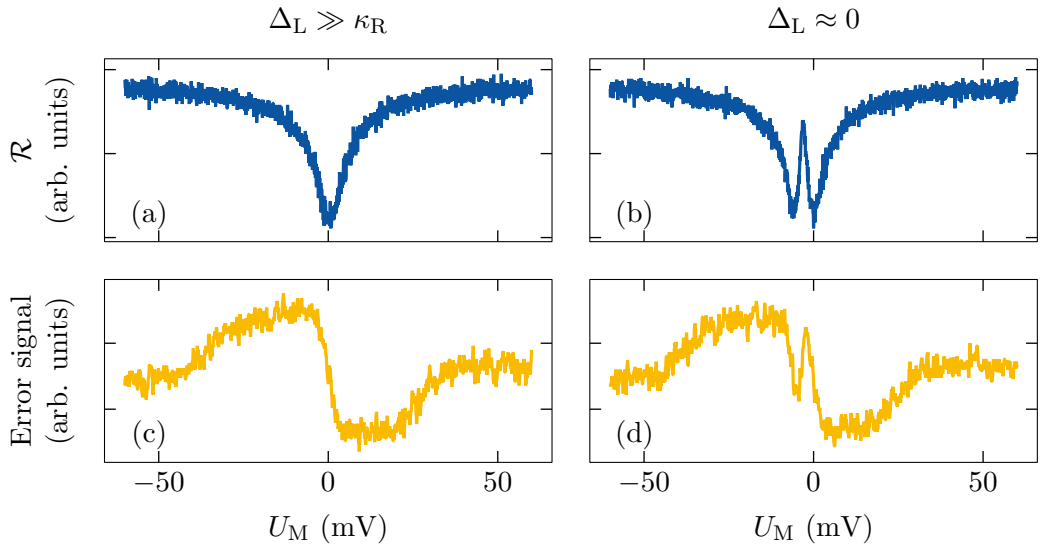


Figure 5.9.: Reflection spectrum of the right cavity with the left cavity being (a) resonant ($\Delta_L = 0$), and (b) far detuned ($|\Delta_L| \gg \kappa_R$). (c) & (d) show the corresponding error signals obtained from the reflected traces, representing the acquired phase when reflecting from the cavity in the respective configuration.

for several values of $\Delta_L = \omega_L - \omega_R$. The fits (red) use Equation (4.4) together with the detuning relations above. Near degeneracy, coupling to the secondary mode produces an asymmetric reduction in reflectance. At degeneracy ($\Delta_L \approx 0$, Subfigure (c)), the two mode amplitudes cancel and the coupling vanishes, revealing a narrow transparency window.

The phase response confirms CRIT-like behaviour: Figure 5.9 compares the right-cavity reflection with the left cavity driven on and off resonance, together with the phase extracted by a Pound–Drever–Hall (PDH) heterodyne readout [83, 84]. A sharp dispersive phase feature accompanies the transparency dip — a hallmark of EIT/CRIT [101, 107].

A steep phase slope implies a reduced group velocity for a broadband wave packet. This property is of particular interest for quantum information processing, where such effects can be exploited for optical storage, where light is slowed down, eventually coming to a full halt [107]. If the required dissipation imbalance can be realised intrinsically, the fibre-based platform becomes particularly appealing because the cavity is naturally fibre-coupled, providing direct and robust in- and outcoupling for photonic signals. However, in our present demonstration, we probe the system using monochromatic light while dynamically tuning cavity lengths to emulate imbalance. To truly realise the slow-light effect, the system must intrinsically satisfy $\kappa_R \gg \kappa_L$, so that a broadband wave packet would experience significant delay within a purely optical setup [108]. One way to achieve this in a future implementations could be intentionally introducing asymmetry in sub-cavity lengths or increasing the losses in one cavity.

In summary, by modulating the detuning of the two sub-cavities at unequal rates we

engineer an effective linewidth imbalance and thereby reproduce CRIT-like interference physics in a purely optical system, despite the absence of intrinsic loss asymmetry and without requiring a long-lived metastable state. This demonstrates that the photonic-molecule platform enables controlled access to interference phenomena that are central to engineered dissipation, and thus provides a versatile testbed for exploring open-system dynamics beyond the parameters fixed by the underlying mirror losses [16].

5.6. Summary

In this chapter we showed how an intrinsic misalignment of the optical cavity axes can be exploited to reduce the cavity–cavity coupling strength g in a controlled, discrete way. The effect arises because misalignment lifts the strict orthogonality of transverse modes on either side of the membrane, creating non-zero spatial overlap between modes that would otherwise be uncoupled.

Because the membrane's finite transmittance sets a minimal coupling strength, the weak-coupling regime is inaccessible if one relies on mirror transmittance alone; the selective overlap control introduced here circumvents that limit and reaches weak coupling. Hence, by selectively addressing different transverse modes, this mechanism lets us realise coupling strengths from $2\pi \times 65$ MHz to $2\pi \times 2.86$ GHz, spanning over almost three orders of magnitude in cooperativity.

We identified the weak-coupling limit via the dissipative interaction of the two cavities, where the secondary cavity produces subtle, detuning-dependent modifications of the probed cavity's amplitude and linewidth, closely analogous to the atomic Purcell effect [100]. By analysing the two-dimensional reflection spectra, we extracted the spectral properties of the secondary cavity even without direct spectroscopy, which is suppressed by the fibre's mode selectivity. Moreover, increasing g revealed the progressive emergence of the minimal mode splitting demonstrating the evolution from a single broadened peak to two well-resolved normal modes, signalling hybridised eigenstates.

Beyond tuning g, we also showed that the effective linewidths can be engineered dynamically, by scanning the two sub-cavities with unequal amplitudes. Consequently, the resonance detuning evolve at different rates, which mimics unequal dissipation ($\kappa_\mathrm{R} \neq \kappa_\mathrm{L}$). This emulation enabled observation of coupled-resonator–induced transparency (CRIT) [14], the purely optical analogue of EIT [101], despite the absence of intrinsic long-lived states. The transparency coincides with a sharp dispersive phase feature measured with a heterodyne measurement technique, pointing towards slow-light in a fibre-coupled platform. However, to take advantage of the effect, we require an intrinsic loss asymmetry [108]

Taken together, the photonic-molecule framework is underpinned by the chapter's observations, showing striking similarities to (linear) atomic phenomena, like Purcell enhancement, mode splitting, and EIT-like transparency [15, 16]. These results establish our device as a tuneable, open photonic molecule: two coupled optical modes whose coherent exchange and effective dissipation can be engineered to model spectra, lifetimes, and interference in different regimes.

6. Absorption-induced nonlinearities in a coupled-cavity environment

Contents

6.1. Characterisation of the nonlinear behaviour in a single cavity resonator 90
6.1.1. Mathematical model 91
6.1.2. Measurements of the bistable optical response 93
6.1.3. Response time of the absorption-induced nonlinearity 98
6.2. Towards unconventional photon blockade effect in coupled cavities 103
6.2.1. Optical response of the nonlinear photonic molecule 104
6.2.2. Optimal parameters for destructive interference 108
6.2.3. Prospects for experimental implementation 111
6.3. Summary 112

In Chapter 3, we identified a pronounced asymmetry in optical losses depending on the orientation of the membrane mirror, attributed to absorption within the silicon nitride layer. This absorption leads to localised heating of the membrane, which in turn, increases its optical thickness through two mechanisms: (1) thermal expansion of the silicon nitride layer [109], and (2) a temperature-induced increase in its refractive index via the thermo-optic effect [110]. Together, these effects give rise to a thermally induced analogue of the optical Kerr effect [111], in which the roundtrip phase acquired by the circulating light becomes intensity-dependent.

Unlike linear optical systems, where the response scales linearly with input power, Kerr-type nonlinearities introduce a nonlinear relationship between input and output fields. In the context of optical cavity resonators, this manifests as optical bistability, a phenomenon in which two distinct intracavity field states (high and low intensity) can exist for the same cavity detuning, depending on the system's prior state [59].

In such bistable systems, the intracavity intensity becomes a nonlinear function of the incident power, enabling photon–photon interactions where the presence of one photon can influence the transmission or reflection experienced by subsequent photons. This property can be utilised to implement all-optical switching, where the optical response of the cavity is modulated by the intensity of the input light — either from the same probe beam or from an independent control beam. These mechanisms are of particular relevance for high-speed, all-optical communication networks [111, 112].

Pushing this nonlinearity to the single-photon level allows the development of deterministic single-photon sources [34, 35], which are key components in emerging quantum technologies such as quantum information networks [113] and quantum key distribution protocols [114]. However, achieving conventional photon blockade — where a single photon inhibits the transmission of additional photons — typically requires strong optical nonlinearities, often realised via active intracavity media like trapped atoms or quantum dots [31, 115].

An alternative approach is offered by the coupled-cavity configuration, where the interaction between a linear cavity and a secondary, weakly nonlinear cavity can enhance the switching behaviour substantially [116]. With optimally tuned parameters, this gives rise to unconventional photon blockade (UPB) through quantum interference effects, where destructive interference prohibits simultaneous coupling of two photons [37, 117]. Remarkably, this mechanism can produce effective photon antibunching even when the nonlinearity is much smaller than the cavity linewidth of the nonlinear cavity [34].

The primary aim of this chapter is to explore how coupling a linear cavity to a weakly nonlinear auxiliary cavity enables an interference-based route to photon-blockade-like behaviour (unconventional photon blockade). In this setting, the coupled system can also be viewed as a nonlinear photonic molecule, i.e. an intensity-dependent extension of the photonic-molecule picture introduced in the previous chapter.

We first characterise the nonlinear response in the single-cavity configuration by directly probing the right cavity that contains the absorptive silicon nitride layer, while keeping the left cavity off resonance. We then investigate the coupled-cavity configuration and analyse how the nonlinear phase response of the right cavity feeds back onto the effective resonance condition of the primarily driven left cavity. Building on the resulting nonlinear dispersion, we discuss conceptually how suitable choices of detuning, coupling, and nonlinearity enable destructive interference between two-photon excitation pathways. Finally, we derive the corresponding optimal parameters within the established semi-classical framework and outline practical modifications that could bring a future experimental demonstration of the unconventional photon blockade within reach.

6.1. Characterisation of the nonlinear behaviour in a single cavity resonator

To model the nonlinear response of the system, we start focusing on the cavity that incorporates the silicon nitride (Si_3N_4) membrane — for consistency referred to as the right (sub-)cavity resonator, with all its defining parameters denoted with the subscript R. As discussed earlier, the absorption of intracavity light by the Si_3N_4 layer leads to localised heating, which in turn modifies the membrane's optical thickness. The resulting thermal shift in resonance frequency introduces an effective intensity-dependent detuning, analogous in form to the optical Kerr effect. In contrast to quasi-instantaneous Kerr nonlinearities, however, this thermally induced shift exhibits a slower dynamical response — an important distinction we revisit at the end of the section. Here, we first develop a mathematical framework that captures the essential features of this nonlinear

behaviour by modelling it as a Kerr-type nonlinearity in the system Hamiltonian. This allows us to analyse the resulting intensity-dependent resonance shift, bistability, hysteresis, and asymmetric lineshape, providing a foundation for understanding the more complex dynamics in the coupled-cavity configuration explored later in this chapter.

6.1.1. Mathematical model

A Kerr-type nonlinearity makes the cavity resonance depend on the intracavity photon number. We describe this by a Kerr coefficient K (shift of the cavity resonance per photon), so that the nonlinear energy is $\hbar K$ per photon pair and hence scales as $(\hat{a}_\mathrm{R}^\dagger \hat{a}_\mathrm{R})^2$, where $\hat{n}_\mathrm{R} = \hat{a}_\mathrm{R}^\dagger \hat{a}_\mathrm{R}$ is the photon-number operator [94]. The Hamiltonian of a driven nonlinear cavity with resonance frequency ω_R is then [116, 117]:

$$\hat{\mathcal{H}} = \underbrace{\hbar\omega_\mathrm{R}\hat{a}_\mathrm{R}^\dagger\hat{a}_\mathrm{R}}_{\text{Linear term}} + \underbrace{\hbar K\left(\hat{a}_\mathrm{R}^\dagger\hat{a}_\mathrm{R}\right)^2}_{\text{Nonlinear term}} + \underbrace{i\hbar F\left(\mathrm{e}^{-i\omega_\mathrm{p}t}\hat{a}_\mathrm{R}^\dagger + \mathrm{e}^{+i\omega_\mathrm{p}t}\hat{a}_\mathrm{R}\right)}_{\text{Pump term}}, \tag{6.1}$$

where F is the pump rate of the incident field with frequency ω_p.

Analogous to the procedure outlined in Section 4.1, from this Hamiltonian, the equation of motion for the mean amplitude of the intracavity field, $\langle\hat{a}_\mathrm{R}\rangle = \alpha_\mathrm{R}$, can be derived as [116]:

$$\begin{aligned}\dot{\alpha}_\mathrm{R} &= \left(i\Delta_\mathrm{p} - \frac{\kappa_\mathrm{R}}{2}\right)\alpha_\mathrm{R} - iK\left\langle\hat{a}_\mathrm{R}^\dagger\hat{a}_\mathrm{R}\hat{a}_\mathrm{R} + \hat{a}_\mathrm{R}\hat{a}_\mathrm{R}^\dagger\hat{a}_\mathrm{R}\right\rangle + F \\ &\approx \left(i\left(\Delta_\mathrm{p} - 2K|\alpha_\mathrm{R}|^2\right) - \frac{\kappa_\mathrm{R}}{2}\right)\alpha_\mathrm{R} + F,\end{aligned} \tag{6.2}$$

where $\Delta_\mathrm{p} = \omega_\mathrm{p} - \omega_\mathrm{R}$ is the detuning between the frequency of the driving field and the cavity resonance, and κ_R is the cavity linewidth, describing dissipation. This approximation assumes the mean-field limit simplifying the nonlinear term $\langle\hat{a}_\mathrm{R}^\dagger\hat{a}_\mathrm{R}\hat{a}_\mathrm{R} + \hat{a}_\mathrm{R}\hat{a}_\mathrm{R}^\dagger\hat{a}_\mathrm{R}\rangle$ as $2\langle\hat{a}_\mathrm{R}^\dagger\hat{a}_\mathrm{R}\rangle\langle\hat{a}_\mathrm{R}\rangle$, which is valid in the regime of large photon numbers, where single photon effects are negligible [118, 119].

In the steady state, the normalised intracavity intensity $|\tilde{\alpha}_\mathrm{R}|^2$, rescaled such that $\max(|\tilde{\alpha}_\mathrm{R}|^2) = 1$, is then given by:

$$|\tilde{\alpha}_\mathrm{R}|^2 = \frac{(\kappa_\mathrm{R}/2)^2}{(\Delta_\mathrm{p} - K_\mathrm{eff}|\tilde{\alpha}_\mathrm{R}|^2)^2 + (\kappa_\mathrm{R}/2)^2}, \tag{6.3}$$

where $K_\mathrm{eff} = (2K)\cdot(2F/\kappa_\mathrm{R})^2$ is the effective nonlinearity, now dependent on the pump strength F^2, and thus on the incident power P_in. The response remains Lorentzian in structure, but the detuning becomes intensity dependent, yielding a self-consistent deformation of the linear resonance.

Solving for the self-consistent response, $|\tilde{\alpha}_\mathrm{R}|^2$, involves finding the (real-valued) roots

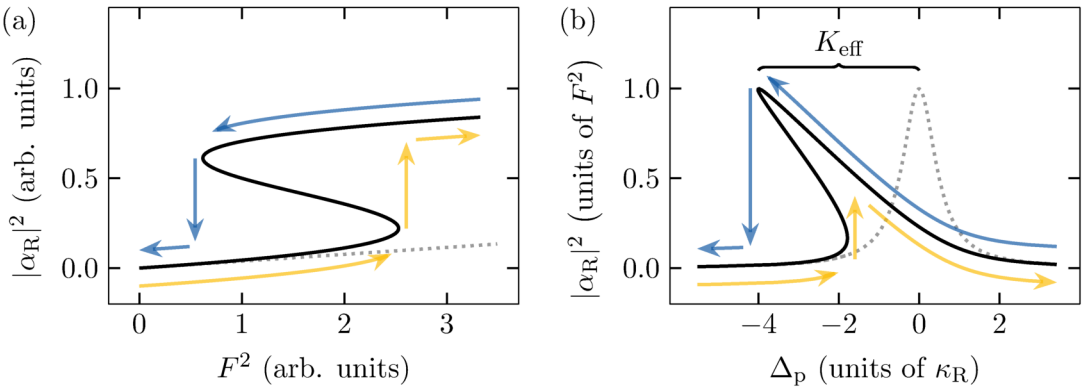


Figure 6.1.: Calculated steady-state intracavity intensities $|\alpha_R|^2$ for a Kerr-nonlinear cavity (solid) and the corresponding linear response ($K = 0$, dotted), shown as a function of (a) pump strength F^2 and (b) probe detuning Δ_p. The curves are evaluated for $(K/\kappa_R) = -2$, with (a) $(\Delta_p/\kappa_R) = -2.5$ fixed and (b) $(F/\kappa_R) = 0.5$ fixed. The coloured arrows indicate adiabatic sweeps of the control parameter, illustrating hysteresis: within the bistable interval (bounded by the vertical arrows), the system follows either a high-intensity (blue) or low-intensity (orange) branch depending on scan direction. Subfigure (b) highlights the resulting asymmetric (skewed) resonance profile, parameterised by the effective nonlinear shift K_{eff}.

of the resulting cubic equation:

$$K_{eff}^2 \left(|\tilde{\alpha}_R|^2\right)^3 - 2K_{eff}\Delta_p \left(|\tilde{\alpha}_R|^2\right)^2 + \left(\Delta_p^2 + \left(\frac{\kappa_R}{2}\right)^2\right)\left(|\tilde{\alpha}_R|^2\right)^1 - \left(\frac{\kappa_R}{2}\right)^2 = 0. \tag{6.4}$$

This cubic equation supports up to three real-valued steady-state solutions for the intracavity intensity. This behaviour is illustrated in Figure 6.1, sketching the non-normalised steady-state solutions for $|\alpha_R|^2$ of a nonlinear cavity (solid curves) together with the corresponding linear response ($K = 0$, dotted curves). In both representations — modulating the pump strength at fixed detuning (Figure 6.1.(a)) or changing the detuning at fixed pump (Figure 6.1.(b)) — there exists a region in the parameter space where multiple solutions coexist. Within this bistable region, two branches are experimentally accessible under adiabatic tuning (a low-intensity and a high-intensity state), while the intermediate branch is not observed in steady scans. Consequently, the observed state depends on the scan history (i.e. how the operating point is reached): under adiabatic scans the intracavity field follows one accessible branch until a turning point is reached, where it switches discontinuously to the other branch [59]. This behaviour is highlighted by the coloured trajectories in Figure 6.1, indicating the resulting hysteresis cycles for opposite scan directions.

Figure 6.1 also explains how the Kerr-type shift shapes the optical response. At finite intensity of the incident light field, the effective detuning becomes intensity dependent ($\Delta_p \to \Delta_p - K_{eff}|\tilde{\alpha}_R|^2$), which skews the Lorentzian response and produces an asymmetric lineshape, as shown in Figure 6.1.(b). In the bistable region, this skewing can

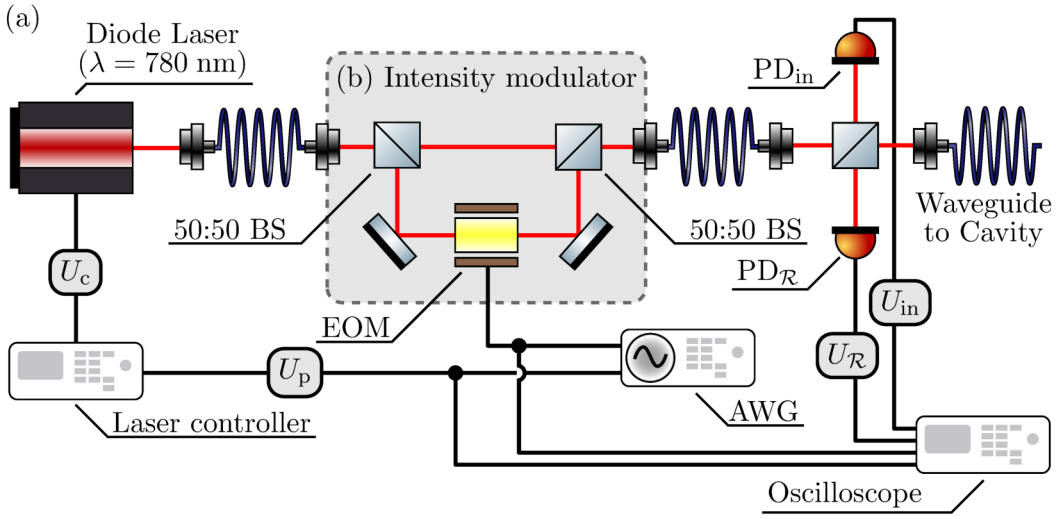


Figure 6.2.: (a) Setup for adiabatic sweep of the input power using an arbitrary waveform generator (AWG) and an interferometric, electro-optic modulator (EOM)-based intensity modulator. An additional photodiode PD_{in} records the incident power in parallel to the cavity interaction. (b) Working principle of the intensity modulator: a voltage-controlled phase shift in one interferometer arm is converted into amplitude modulation by interference at the recombining beam splitter. This figure includes elements generated with Ref. [54].

be understood geometrically: the resonance shift increases with intensity, so the high-intensity solution is displaced in detuning relative to the low-intensity solution, and the resulting deformation "wraps" one flank of the response over the other. As a result, the low- or a high-intensity steady state can be associated with the red- and blue-detuned sides of the cavity resonance (with the assignment depending on the sign of K).

6.1.2. Measurements of the bistable optical response

With this picture in mind, demonstrating and analysing optical bistability requires systematic control of both the laser detuning Δ_p and the incident power P_{in}. To this end, we modify the optical setup introduced in Figure 4.8 in two key ways. First, we replaced the electro-optic modulator (EOM) previously used to generate frequency markers by a fibre-coupled intensity modulator[1]. By converting an applied voltage ramp into a controlled power ramp, this device enables adiabatic tuning of the incident intensity and thereby access to both branches of the bistable response, even for constant detuning. Second, we added an additional photodiode to record the incident and reflected powers simultaneously, yielding the voltages U_{in} and $U_{\mathcal{R}}$, respectively. A schematic of the modified setup, including the operating principle of the intensity modulator, is shown in Figure 6.2.

This configuration enables a two-dimensional mapping of the optical response as a

[1] iXblue: NIR-MX800-LN-20-00-P-P-FA-FA

function of Δ_p and P_{in} by combining a fast detuning scan with a slow power ramp, analogous to the scheme in Section 4.3.1. The probe frequency ω_p is scanned via the laser controller driven by a triangular waveform U_p from an arbitrary waveform generator (AWG) at $f_p = 50$ Hz, while the incident power is ramped over a time window of 3500 ms using the intensity modulator. Each data point of the reflected signal can thus be assigned a detuning and an input power value, enabling construction of a two-dimensional dataset.

The horizontal axis (detuning Δ_p) is calibrated in frequency units by fitting Lorentzian profiles to horizontal cross-sections and using the independently determined linewidth $\kappa_R = 168(3)$ MHz (Section 4.2). For this calibration, the probe power is kept sufficiently low to avoid nonlinear effects; the observation of an approximately constant linewidth at decreasing input powers is consistent with negligible nonlinearity in this regime. The vertical axis (voltage U_{in}) is converted to absolute incident power using the calculated photodiode responsivity [2].

The resulting reflection maps acquired during a single incident-power ramp are shown in Figure 6.3.(a)–(b). Rather than plotting the raw reflected power $P_{\mathcal{R}}$, we evaluate the difference with respect to an off-resonant reference measurement:

$$\Delta P_{\mathcal{R}}(\Delta_p) = P_{\mathcal{R}}(\Delta_p \to \infty) - P_{\mathcal{R}}(\Delta_p), \tag{6.5}$$

so that $\Delta P_{\mathcal{R}}$ directly tracks the power coupled into the cavity and is proportional to the normalised intracavity intensity $|\tilde{\alpha}_R|^2$. This differential representation also suppresses the trivial background increase of $P_{\mathcal{R}}$ with rising P_{in}, thereby improving the contrast of the resonance.

The spectra shown in Figure 6.3.(a)–(b) reveal different optical responses for opposite scan directions.This asymmetry is a clear signature of optical hysteresis, with the bistable region given by the parameter area between the contours of maximal response traced for the two opposite scan directions in the (Δ_p, P_{in}) plane, as indicated by the dashed white lines. For $\dot{\Delta}_p < 0$, the trajectory follows the branch that reaches the high-intensity state before switching, characterised by stronger intracavity fields, whereas for $\dot{\Delta}_p > 0$ the system follows the complementary low-intensity branch. This behaviour is consistent with the qualitative picture introduced in Section 6.1.1.

To further examine the bistable response, horizontal and vertical cross-sections of the two-dimensional reflection spectra are presented in Figure 6.3.(c)–(d)[3]. These plots show the optical response at fixed detuning $\Delta_p = -2\pi \times 411$ MHz and fixed input power $P_{in} = 72(1)$ µW, as indicated by the red lines in Figure 6.3.(a)–(b). The high- and low-intensity branches are shown in blue and orange, respectively, in accordance with the colours of the trajectories in Figure 6.1. The black curve shows the corresponding steady-state solution for the appropriately scaled $|\tilde{\alpha}_R|^2$ obtained from Equation (6.4).

The good agreement between measurement and model confirms that, under adiabatic

[2]Thorlabs: PDA36A2: Si Switchable Gain Detector

[3]Figure 6.3.(d) is the horizontal cut taken directly along the solid lines in (a) and (b). Figure 6.3.(c) was recorded in a separate fixed-detuning scan (using a wavelength meter) to resolve both switching transitions with sufficient resolution; its position is therefore indicated only approximately (dotted line).

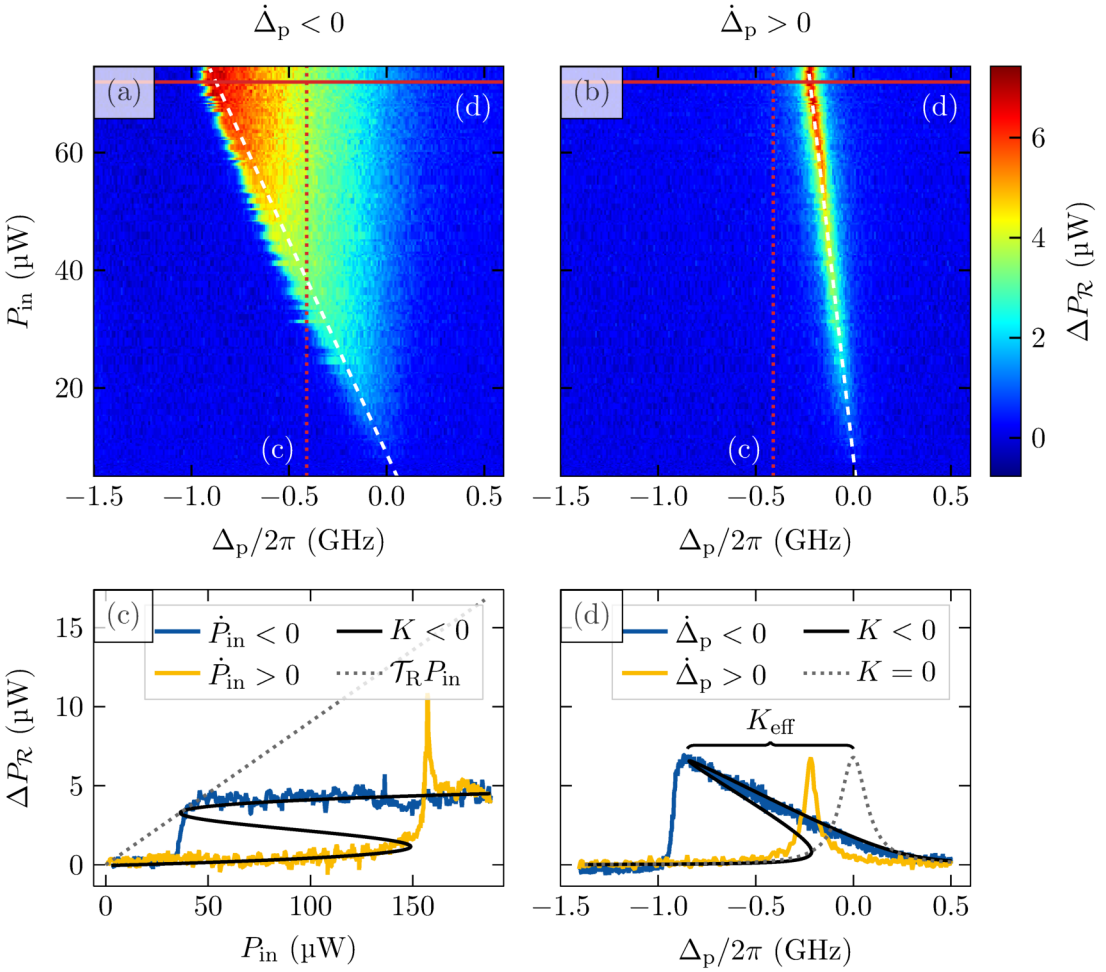


Figure 6.3.: (a) & (b) Two-dimensional reflection spectra of the cavity, recorded while scanning the probe detuning Δ_p (horizontal axis) and ramping the incident power P_in (vertical axis). The two subfigures correspond to opposite scan directions of the laser-frequency, revealing optical hysteresis. The white dashed curves trace the detuning of maximal coupled power $\Delta P_\mathcal{R}$ as a function of input power and define the bistable region. (c) & (d) Cross-sections illustrating the coupled power for the low- and high-intensity states, together with the corresponding (appropriately scaled) theoretical nonlinear response for $K < 0$ (solid black), exhibiting sharp transitions at the turning points of the bistable solution. Subfigure (c) shows $\Delta P_\mathcal{R}$ versus P_in at fixed detuning $\Delta_\mathrm{p} = -2\pi \times 411$ MHz; the dotted line indicates the maximal coupling attainable on resonance. Subfigure (d) shows $\Delta P_\mathcal{R}$ versus Δ_p at fixed power $P_\mathrm{in} = 72(1)$ µW (solid red line in (a) & (b)), and includes the linear response ($K = 0$, dotted). The trace in (d) is extracted from the horizontal profile of (a) & (b), whereas (c) is obtained from a separate measurement (indicated by the red dotted line) to better resolve the power dependence and both switching transitions.

parameter tuning, the intracavity field follows the predicted bistable response. The sharp transitions occur at the points of the steady-state solution, where the followed branch turns and the system switches to the alternative accessible branch. As the two branches are associated with opposite (red- and blue-detuned) flanks of the resonance, a transition corresponds to a jump between these flanks therefore crossing the resonance, where coupling is maximal. This is particularly visible in Figure 6.3.(c), where the low-to-high transition exhibits a pronounced peak that approaches maximal coupling (dotted line), that is not described by the non-dynamic steady-state model.

Extraction of the photothermal coefficient

We can utilise the high-to-low transition to extract the effective nonlinearity because it coincides with the system passing through resonance, producing maximal coupling. As the detuning acquires an intensity-dependent shift, the effective resonance condition is reached when $\Delta_\mathrm{p} - K_\mathrm{eff}|\tilde{\alpha}_\mathrm{R}|^2 = 0$. Since the high-intensity branch reaches its maximum at $|\tilde{\alpha}_\mathrm{R}|^2 \approx 1$, the detuning at the transition directly yields K_eff (see Figure 6.3.(d)).

Repeating this extraction for a range of incident powers P_in (with $P_\mathrm{in} \propto F^2$, where F denotes the pump rate) yields the power dependence of the effective nonlinear shift and thus provides insight into the underlying nonlinear behaviour. We parameterise this dependence as:

$$K_\mathrm{eff} = \beta P_\mathrm{in}, \tag{6.6}$$

where the photothermal coefficient β converts incident power into an effective detuning. This linear dependence is verified in Figure 6.4, which plots the detuning position extracted from the high-intensity branch (represented by maximal $\Delta P_\mathcal{R}$) as a function of P_in using the data from Figure 6.3.(a). A linear fit (red) yields $\beta = -2\pi \times 13.8(2)$ MHz/mW. To account for systematic uncertainty in the zero-detuning reference, the fit also includes an offset.

The sign of β confirms that increasing optical power induces a redshift of the cavity resonance, consistent with a lengthening of the optical path due to thermal expansion of the silicon nitride layer. A frequency shift of the cavity resonance, $\Delta\omega_\mathrm{R}$, can be related to a change in the optical cavity length L_R, via:

$$\Delta\omega_\mathrm{R} = -\frac{2\pi c}{\lambda}\frac{\Delta L_\mathrm{R}}{L_\mathrm{R}}, \tag{6.7}$$

with $L_\mathrm{R} \approx n_\mathrm{SiN} \cdot t_\mathrm{SiN} + d_\mathrm{R}$, includes both the silicon nitride layer and the length of the right sub-cavity. Here, n_SiN and t_SiN denote the refractive index and thickness of the silicon nitride film, while d_R is the geometric cavity length. This model assumes that the silicon nitride effectively acts as an intracavity medium — an assumption validated by transfer-matrix simulations, which show that intracavity intensity and intensity within the silicon nitride layer are proportional (see Appendix A.2).

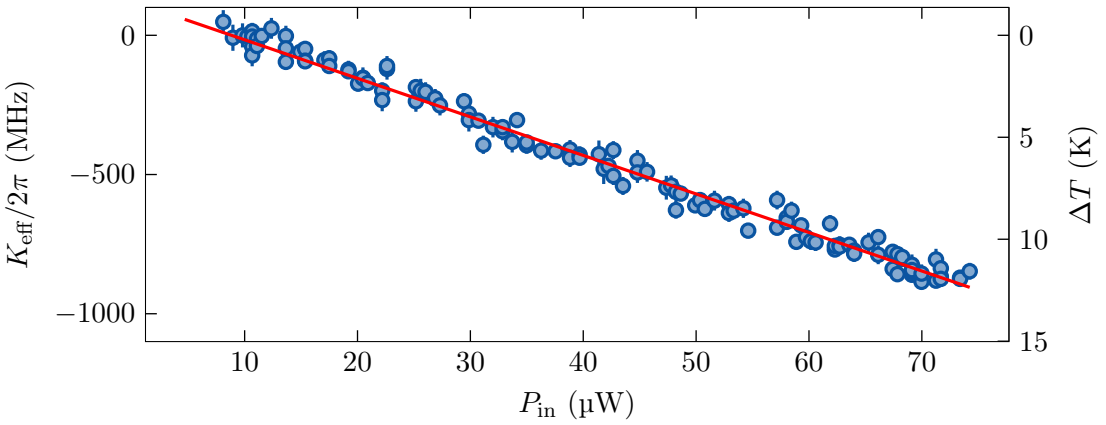


Figure 6.4.: Effective nonlinear detuning K_{eff} extracted from the bistable high-intensity branch as a function of incident power P_{in}. A linear fit (red) yields the photothermal coefficient $\beta = -2\pi \times 13.8(2)$ MHz/mW. The secondary vertical axis indicates the corresponding temperature increase ΔT, inferred from the cavity geometry and the thermal properties of silicon nitride using Equation (6.9).

The temperature-induced change in optical cavity length can be approximated as:

$$\begin{aligned}\Delta L_{\mathrm{R}} &\approx \left(n_{\mathrm{SiN}}\frac{\partial t_{\mathrm{SiN}}}{\partial T} + t_{\mathrm{SiN}}\frac{\partial n_{\mathrm{SiN}}}{\partial T} + \frac{\partial d_{\mathrm{R}}}{\partial T}\right)\Delta T \\ &= \left((n_{\mathrm{SiN}}-1)\alpha_d + \alpha_n\right) t_{\mathrm{SiN}}\Delta T,\end{aligned} \tag{6.8}$$

where α_d and α_n are the coefficients of thermal expansion [109] and thermo-optic response [110], respectively. The term $(n_{\mathrm{SiN}} - 1)$ accounts for the displacement of air as the silicon nitride layer expands, following $\partial_T d_{\mathrm{R}} = -\partial_T t_{\mathrm{SiN}}$.

Combining Equation (6.7) and Equation (6.8), we obtain the temperature increase ΔT as a function of the observed resonance shift $\Delta\omega_{\mathrm{R}} = K_{\mathrm{eff}} = \beta P_{\mathrm{in}}$:

$$\Delta T = \underbrace{\frac{\beta\lambda}{2\pi c\left((n_{\mathrm{SiN}}-1)\alpha_d + \alpha_n\right)}\frac{L_{\mathrm{R}}}{d_{\mathrm{SiN}}}}_{\approx 200\ \frac{\mathrm{mK}}{\mu\mathrm{W}}} \cdot P_{\mathrm{in}}. \tag{6.9}$$

Previous studies have shown a linear relationship between α_n and the refractive index n_{SiN} [110]. Although their measurements were performed at $\lambda = 1550$ nm, we may expect that the scaling is primarily dependent on the (also wavelength-dependent) refractive index, and not significantly on the absolute wavelength value itself. Using our measured refractive index $n_{\mathrm{SiN}} \approx 2.17(1)$ (see Section 3.2.1) and interpolated values from their work, we estimate $\alpha_n \approx 5.5 \times 10^{-5}\ \mathrm{K}^{-1}$. Together with $\alpha_d \approx 3.27 \times 10^{-6}\ \mathrm{K}^{-1}$ [109], this model predicts a maximal temperature increase of approximately 200 mK/µW in response to the incident light field. The estimated temperature change is also shown as

secondary vertical axis in Figure 6.4.

This analysis highlights that the supposedly temperature-induced shift of the cavity resonance is dominated by the thermo-optic effect, rather than by thermal expansion. Despite the small absolute temperature changes, the cumulative effect on the optical path — due to the thickness of the silicon nitride layer — can induce resonance shifts of several cavity linewidths.

Effective nonlinear refractive index

To put the strength of the nonlinearity in perspective, we can estimate the underlying change in the refractive index, analogously to the Kerr effect [111], where the refractive index is modulated by an additional intensity-dependent term:

$$\begin{aligned} n'_{\mathrm{SiN}} &= n_{\mathrm{SiN}} + n_2 \cdot I_{\mathrm{circ}} \\ &\approx n_{\mathrm{SiN}} + n_2 \cdot \frac{1}{2}\left(\frac{\pi w_0^2}{2}\right)\left(\frac{\mathcal{F}_{\mathrm{R}}}{\pi} \cdot P_{\mathrm{in}}\right), \end{aligned} \tag{6.10}$$

where n_2 denotes the nonlinear component of the refractive index and w_0 is the waist of the cavity mode, and the additional factor of $1/2$ results from the temporal average of the oscillating standing wave pattern inside the cavity resonator. Finally, the factor $(\mathcal{F}_{\mathrm{R}}/\pi)$ relates the input power to the power circulating inside the cavity, taking its resonant amplification into account. Assuming that the shift in resonance, described in Equation (6.7), is only caused by a change in the silicon nitride's refractive index, $\Delta n = n_2 \cdot I_{\mathrm{circ}}$, according to $\Delta L_{\mathrm{R}} \approx \Delta n \cdot t_{\mathrm{Sin}}$, then reveals an equivalent nonlinear refractive index component of:

$$n_2 \approx 2.6 \times 10^{-16}\ \frac{\mathrm{m}^2}{\mathrm{W}}. \tag{6.11}$$

This value is orders of magnitude larger than conventional nonlinear dielectric media [120–122], where their nonlinearity is caused by purely optical effects.

6.1.3. Response time of the absorption-induced nonlinearity

The photothermal nonlinearity is expected to depend on the duration over which the membrane is illuminated: faster scans reduce the effective heating time, while slower scans allow the response to build up toward a steady state. We therefore study how the extracted nonlinearity parameter β depends on the absolute scan rate, $|\dot{\Delta}_{\mathrm{p}}|$. The resulting relationship both tests the thermal hypothesis and enables an order-of-magnitude estimate of the nonlinear response time.

To this end, we measure the effective nonlinearity for different scan rates $|\dot{\Delta}_{\mathrm{p}}|$ and incident powers P_{in}. In order to retain measurable nonlinear shifts even at fast scans, we remove the intensity modulator, whose insertion loss and power-handling limit constrained the available P_{in}. Instead, we employ a frequency sweep that increases the scan rate dynamically during the acquisition. Repeating this procedure for different input

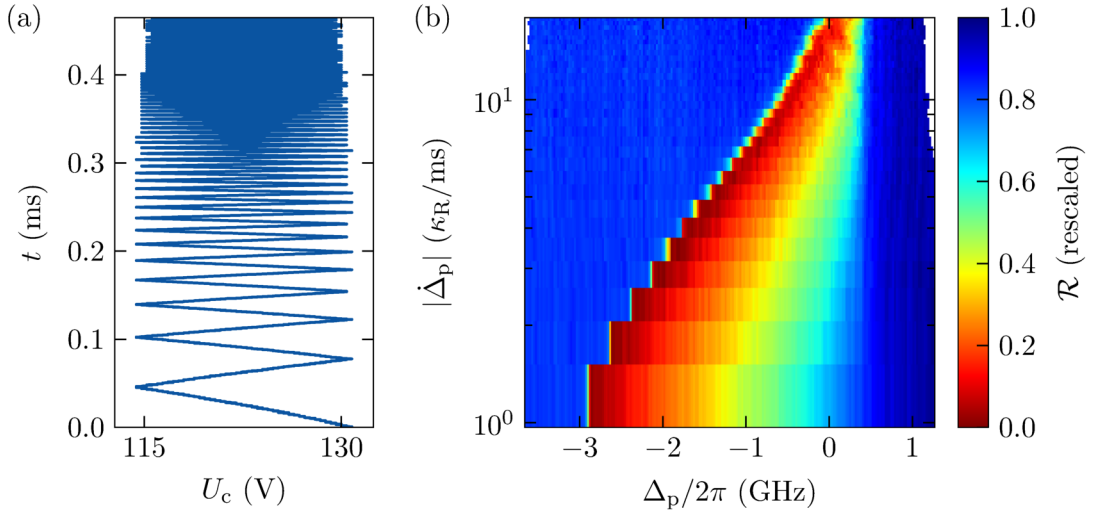


Figure 6.5.: (a) Output voltage U_c measured at the laser controller output, showing the applied frequency-sweep waveform with increasing scan rate at later times t. (b) Reflection spectrum recorded for the high-intensity (bistable) branch using the sweep in (a) at $P_{in} = 275(13)$ µW. The horizontal axis gives the reflected signal amplitude, while the vertical axis encodes the absolute value of the scan rate $|\dot{\Delta}_p|$ extracted from the local slope of U_c. Both the detuning scale and $\dot{\Delta}_p$ are calibrated using the cavity linewidth as reference. At large scan rates, the reduced scan range reflects the finite bandwidth of the laser controller.

powers yields data points of the form $(P_{in}, K_{eff}, |\dot{\Delta}_p|)$. By sorting them into narrow bins of comparable $|\dot{\Delta}_p|$ and analysing the power dependence within each bin, we extract the photothermal coefficient β as a function of $|\dot{\Delta}_p|$.

An exemplary frequency-sweep measurement is shown in Figure 6.5 for an input power of $P_{in} = 275(13)$ µW. Figure 6.5.(a) displays the measured controller output U_c, illustrating the increasing scan frequency (and thus $|\dot{\Delta}_p|$) over the course of the acquisition. We use the controller output as in practice, the finite bandwidth of the laser controller used to interface with the laser attenuates fast components and introduces phase delay at high scan rates. The direct use of U_c to map the laser detuning, instead of the programmed AWG signal U_p, circumvents this limitation. Therefore a junction is implemented between controller and laser, allowing simultaneous monitoring of both U_c and U_p. The corresponding pick-off points are indicated in Figure 6.2.

Figure 6.5.(b) shows the corresponding reflection spectrum for the scan direction accessing the high-intensity branch ($\dot{\Delta}_p < 0$). The data are plotted as horizontal line profiles versus detuning Δ_p. The absolute scan rate $|\dot{\Delta}_p|$, obtained experimentally from the slope of each individual ramp segment of U_c, is shown on the vertical axis. The conversion from U_c to detuning is again calibrated using the independently measured cavity linewidth $\kappa_R = 168(3)$ MHz as a reference. In the limit of low optical power and sufficiently fast scans, the response approaches the linear cavity lineshape, providing a practical calibration point for the detuning scale; these conditions were verified experi-

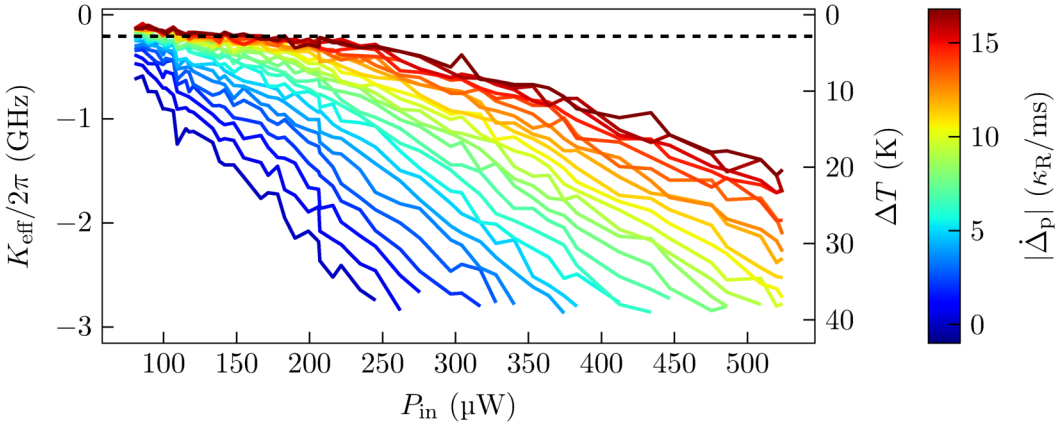


Figure 6.6.: Effective nonlinear detuning K_{eff} extracted from frequency-sweep measurements as a function of incident power P_{in} for a set of scan rates $\dot{\Delta}_{\mathrm{p}}$. Each trace combines multiple sweeps grouped by P_{in} and colour-coded by $|\dot{\Delta}_{\mathrm{p}}|$. The right axis indicates the corresponding temperature change ΔT. The dashed line marks the approximate onset of bistability; below this threshold the response remains single-branched, so an effective linewidth of the high-intensity branch is no longer defined in the same way, which alters the relationship between P_{in} and K_{eff}.

mentally. For better context, we express the scan rate in units of (κ_{R}/ms), corresponding to the number of cavity linewidths scanned per millisecond.

The reconstructed power dependence of the effective nonlinearity is shown in Figure 6.6. Here, K_{eff} is again extracted from the detuning at which the high-intensity branch reaches maximal coupling and transitions back to the low-intensity state. Overall, the data exhibit an approximately linear scaling with input power, with a slope that increases for slower scan rates (represented by the colour scale), consistent with the gradual build-up of the photothermal resonance shift.

At high scan rates and low input powers, the slope decreases noticeably, indicating a change in scaling. In this regime, the nonlinear shift is too weak to establish bistability, so the response remains single-branched and the effective heating time is set primarily by the scan itself. Once bistability emerges, increasing P_{in} not only enhances the resonance shift but also broadens the coupling window of the high-intensity branch, thereby extending the time over which the membrane is pumped and heated. This crossover between the regimes coincides with the dashed line in Figure 6.6, marking the threshold for bistable behaviour.

To obtain a consistent estimate of the photothermal coefficient, we therefore extract β only in the regime where bistability is fully established, by applying a linear fit to the high-power data well above the threshold. The resulting values of $\beta(|\dot{\Delta}_{\mathrm{p}}|)$ are summarised in Figure 6.7 and compared with the value obtained from the power-ramp measurement in Section 6.1.2.

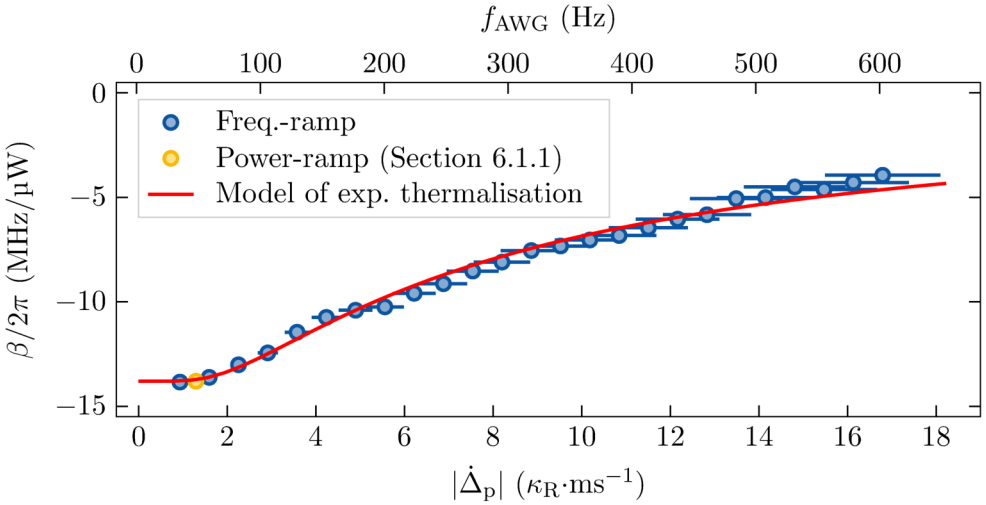


Figure 6.7.: Effective photothermal coefficient β as a function of the absolute scan rate $|\dot{\Delta}_{\mathrm{p}}|$ of the cavity detuning. Frequency-sweep measurements (blue) are compared to the power-ramp result (orange; see Section 6.1.2) and to an exponential-thermalisation model for the silicon nitride layer (red). For ease of comparison between measurement protocols, the upper axis converts $|\dot{\Delta}_{\mathrm{p}}|$ into the AWG scan frequency f_{AWG}, assuming a detuning sweep range of $2\pi \times$ 2.5 GHz, consistent with Figure 6.3.

Model of exponential thermalisation

The scan-rate dependence of the extracted photothermal coefficient $\beta(|\dot{\Delta}_{\mathrm{p}}|)$ follows the qualitative expectation for a thermal-driven process: at slow scans the membrane can approach a steady-state temperature, whereas at fast scans the optical pumping is too brief for the full temperature rise to develop. A simple description of this behaviour is an exponential approach to steady state, as is commonly used for thermal dynamics [123, 124]. Because β is proportional to the temperature increase of the membrane (see Equation (6.9)), it can be used as an indirect thermometer.

To connect scan rate to heating, we approximate that the membrane is effectively pumped over a finite detuning window during the scan, so that the relevant heating time Δt scales inversely with the absolute scan rate $|\dot{\Delta}_{\mathrm{p}}|$. Under this assumption, the exponential thermalisation model can be written as:

$$\Delta T(\Delta t) = \Delta T_{\mathrm{ss}} \left(1 - \exp\left(-\frac{\Delta t}{\tau_T}\right)\right) \quad \rightarrow \quad \beta(\dot{\Delta}_{\mathrm{p}}) \approx \beta_{\mathrm{ss}} \left(1 - \exp\left(-\frac{\Lambda}{|\dot{\Delta}_{\mathrm{p}}|}\right)\right), \quad (6.12)$$

where β_{ss} is the steady-state photothermal coefficient at slow scan rates. The parameter Λ sets the characteristic scan-rate scale at which the response crosses over from the fast-scan to the steady-state limit, and τ_T denotes the corresponding photothermal response time.

Fitting Equation (6.12) to the data in Figure 6.7 captures the observed crossover well,

supporting the interpretation in terms of exponential heating dynamics. The extracted steady-state value $\beta_{\mathrm{ss}} = -2\pi \times 13.8(1)$ MHz/mW is consistent with the independent estimate obtained from the power-ramp measurement in Section 6.1.2.

The response time τ_T is related to Λ through the characteristic product of absolute scan rate and effective heating time:

$$\tau_T = \frac{\Delta t \cdot |\dot{\Delta}_{\mathrm{p}}|}{\Lambda}, \tag{6.13}$$

as implied by Equation (6.12). In practice, however, Δt is not purely set by the scan rate: increasing input power increases the resonance shift and thereby broadens the detuning interval over which significant intracavity power is present. This introduces a weak self-consistency, such that Δt depends on the very nonlinearity that is being extracted from. For this reason, we restrict ourselves to an order-of-magnitude estimate by taking a characteristic heating window of $\Delta t \cdot |\dot{\Delta}_{\mathrm{p}}| \approx 2\pi \times 1$ GHz[4], yielding:

$$\begin{aligned} \tau_{\mathrm{T}} = \frac{\Delta t \cdot |\dot{\Delta}_{\mathrm{p}}|}{\Lambda} &\approx \frac{2\pi \times 1 \text{ GHz}}{\Lambda} \\ &\approx 25(2) \text{ µs}. \end{aligned} \tag{6.14}$$

Given the high thermal conductivity of Si_3N_4 [125], thermo-optic/photothermal time constants in the order of tens to hundreds of microseconds are plausible in micro-photonic silicon nitride structures, consistent with reported thermal lifetimes in Si_3N_4 resonators [126, 127].

It is instructive to compare this timescale to alternative mechanisms. A potential optomechanical contribution would be expected to evolve on the mechanical relaxation timescale $\tau_{\mathrm{M}} = Q/(2\pi\nu_{\mathrm{M}})$ [128, 129], which for the lowest membrane mode ($\nu_{\mathrm{M}} = 40.5$ kHz and $\mathcal{O}(Q) \approx 10$, see Section 3.4.1) gives $\tau_{\mathrm{M}} \approx 40$ µs, i.e. comparable in order of magnitude. However, the absence of any corresponding bistability signature when probing the left cavity rules out an optomechanical origin in the present system.

By contrast, an electronic Kerr response would be effectively instantaneous on the timescales relevant here ($\mathcal{O}(\tau_{\mathrm{Kerr}}) \sim$ ps [130]) and would therefore not produce the pronounced scan-rate dependence observed in $\beta(\dot{\Delta}_{\mathrm{p}})$. Taken together, the measured scan-rate dependence and the extracted microsecond response time provide strong evidence that the observed nonlinearity is dominated by absorption-induced heating of the silicon nitride layer.

It is worth noting that the above analysis focuses on the high-intensity branch of the bistable response, which dominates the intracavity energy stored during a scan and therefore provides the clearest access to the photothermal response time. The same

[4]The numerical value $2\pi \times 1$ GHz is a representative "effective pumping bandwidth" during a sweep. In Figure 6.6, the extracted nonlinear shifts K_{eff} span roughly $2\pi \times (0.1\text{–}3.0)$ GHz. Taking a typical mid-range shift and estimating the frequency interval over which the intracavity power stays appreciable (of order the effective linewidth/FWHM of the skewed response; approximated here by a simple triangular profile) yields a characteristic window on the order of $2\pi \times 1$ GHz.

reasoning applies to the low-intensity branch, but with a reduced effective heating time because the intracavity power remains small throughout a single ramp. Equivalently, the detuning interval over which the cavity is efficiently pumped is narrower on the low-intensity branch. For example, Figure 6.3.(d) indicates that the coupling window of the low-intensity branch is approximately an order of magnitude smaller than that of the high-intensity branch, such that comparable heating requires a proportionally reduced scan rate. Because this window depends on the incident power through the nonlinear resonance shift, the precise scaling factor is generally power dependent and should be understood as an order-of-magnitude guideline rather than a universal constant.

6.2. Towards unconventional photon blockade effect in coupled cavities

In the previous section we showed that absorption-induced heating of the silicon nitride membrane produces a thermally driven resonance shift of the right cavity, which can be modelled as an effective Kerr-type nonlinearity. In a single nonlinear cavity, this introduces an intensity-dependent modification of the cavity's resonance condition, manifesting as photon-number-selective optical response.

Coupling a second, primarily linear cavity mode to this weakly nonlinear mode allows multi-photon states to be reached via distinct sequences of intermediate states involving both cavities. Because these pathways populate the two sub-cavities differently, they experience different nonlinear resonance shifts and thus acquire different power-dependent phases. As a result, different excitation pathways leading to the same multi-photon state can interfere, reshaping the optical response in close analogy to intensity-independent effects in coupled-resonator systems, such as coupled-resonator-induced transparency [14], as discussed in Section 5.5.

Figure 6.8.(a) illustrates this mechanism by sketching the level structure of the coupled system up to the two-photon manifold. The basis state $|nm\rangle$ denotes the photon number in the primary (left) and auxiliary (right) cavities, respectively, where the auxiliary cavity incorporates the nonlinear silicon nitride layer. Occupation of the auxiliary mode shifts the corresponding energy levels by an amount proportional to K, the single-photon nonlinearity. For a drive applied to the left cavity with rate F, the two-photon state $|20\rangle$ can be reached either directly, $|10\rangle \xrightarrow{F} |20\rangle$, or via the auxiliary cavity, $|10\rangle \xrightarrow{g} |01\rangle \xrightarrow{F} |11\rangle \xrightarrow{g} |20\rangle$. Under suitable conditions these two pathways acquire opposite phases and interfere destructively, suppressing the population of the $|20\rangle$ state. This mechanism is known as the unconventional photon blockade [34] and provides an efficient route to strong antibunching, by allowing only single-photon excitations, even for moderate nonlinearities [37].

The same mechanism can be interpreted in terms of the circulating-field phase, providing a semi-classical description of the effect. As illustrated in Figure 6.8.(b), inter-cavity coupling causes the field in the primary cavity to pick up an additional phase δ upon reflection from the membrane mirror, which depends on the state of the auxiliary cav-

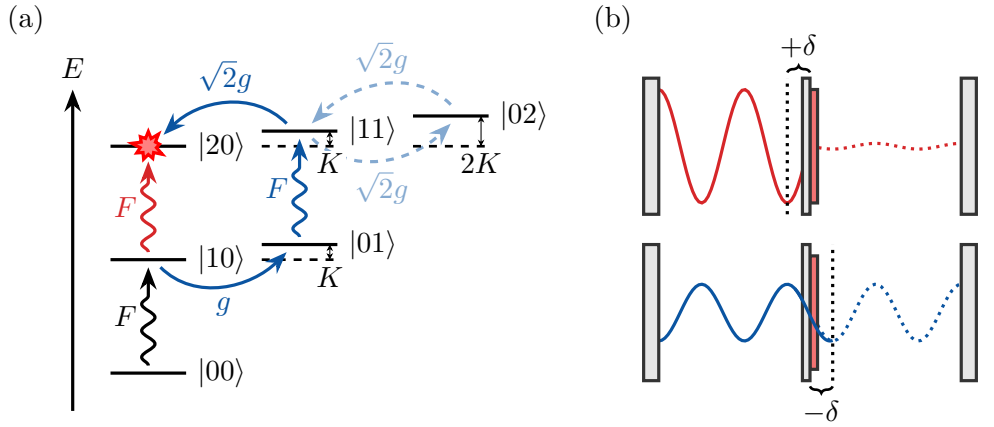


Figure 6.8.: (a) Level scheme of the coupled-cavity system (adapted from Ref. [37]). The basis state $|nm\rangle$ denotes n photons in the left cavity and m photons in the right cavity. Occupation of the nonlinear right cavity shifts the corresponding energies by $m\hbar K$. Red and blue arrows indicate two excitation pathways populating the $|20\rangle$ state. Destructive interference between these pathways suppresses $|20\rangle$, giving rise to photon-blockade-like behaviour. (b) Sketch of the intracavity fields illustrating the same mechanism: the two paths leading to the $|20\rangle$ state acquire phases $\pm\delta$ when reflecting from the membrane mirror, yielding a relative phase difference of π and hence interfere destructively.

ity. By appropriate choice of coupling strength, detuning, and nonlinearity, the relative phase between the two-photon excitation pathways can approach $2\delta = \pi$, leading to destructive interference and thereby inhibiting two-photon excitation.

To connect this semi-classical interference picture to the experimental system and formulate the optimal parameters for destructive interference, we first characterise the power-dependent optical response of the coupled resonators, i.e. the nonlinear photonic molecule formed by coupling the linear cavity to the nonlinear auxiliary cavity. For this, we probe the system through the left cavity, while the right cavity acts as an auxiliary element providing the nonlinear response. With this choice, the measured spectrum remains close to a linear cavity response over most detuning, and signatures of the nonlinearity arise predominantly through coupling-induced back-action in the vicinity of the degeneracy $\omega_\mathrm{L} \approx \omega_\mathrm{R}$. This allows us to isolate coupled-cavity effects that would otherwise mix with single-cavity nonlinearities.

6.2.1. Optical response of the nonlinear photonic molecule

The reflection spectrum of the nonlinear photonic molecule reveals how the inter-cavity coupling g, the cavity–cavity detuning Δ_L, the probe detuning Δ_p, and the effective nonlinearity K_eff shape the coupled-cavity response. As discussed in Section 5.2, the coupled-cavity system can alternatively be described in terms of an effective single-cavity response, where the auxiliary cavity acts as a frequency-dependent reflector. In this picture, the dispersive back-action of the auxiliary cavity is directly linked to the phase

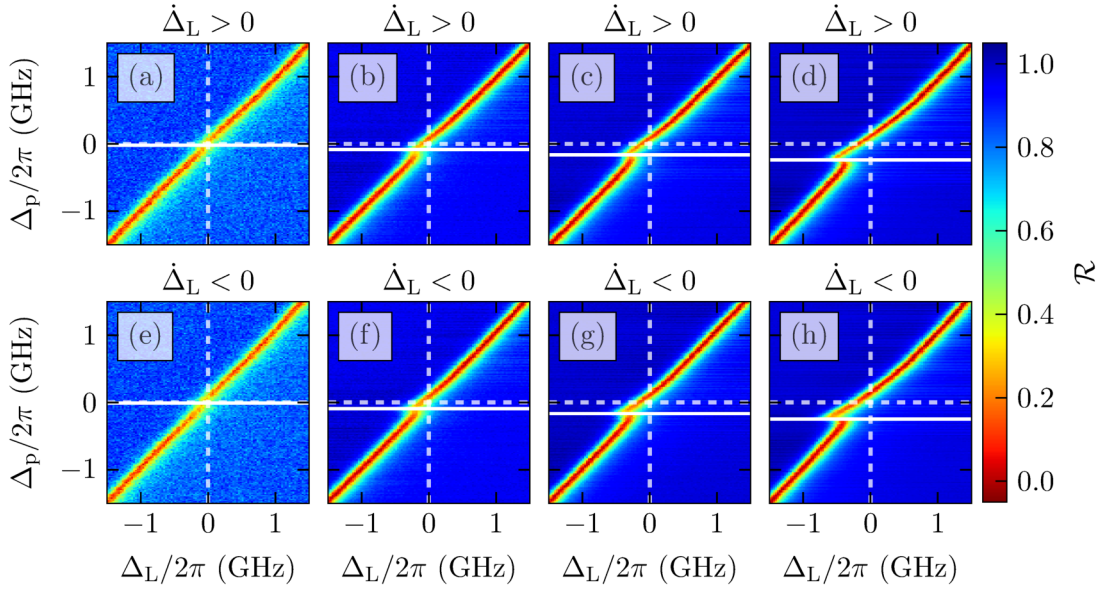


Figure 6.9.: Two-dimensional reflection spectra of the coupled-cavity system in the weak-coupling regime, with coupling strength $g = 2\pi \times 68(5)$ MHz and linewidths $\kappa_L = 2\pi \times 212(2)$ MHz and $\kappa_R = 2\pi \times 212(8)$ MHz. From left to right the incident power is increased, $P_{in} = \{35(2)\ ,\ 205(14)\ ,\ 375(22)\ ,\ 627(34)\}$ µW, revealing the progressive formation of a power-dependent feature caused by the nonlinear response of the auxiliary (right) cavity mode. Subfigures (a)–(d) show spectra acquired for positive scan rates of the left-cavity length modulation, while (e)–(h) show the corresponding spectra for the opposite scan direction. Dashed lines mark the origin of the axes, and the solid horizontal line indicates the probe detuning at which the resonance amplitude is minimal.

of its reflected field. Motivated by this correspondence, we measure reflection spectra as functions of Δ_L and Δ_p for different input powers P_{in} and extract the corresponding centre frequency of the effective single-cavity resonance (see Equation (5.4b)), yielding the power-dependent coupled-system resonance condition.

For this analysis, we focus on the weak-coupling regime. Here, the degeneracy point is fully resolved, allowing the nonlinear shift of the right-cavity resonance and the resulting response of the photonic molecule to be extracted reliably. In the strong-coupling regime the degeneracy is obscured by an avoided crossing (e.g., compare Figure 5.3.(b) to Figure 5.4.(b)), which makes a direct identification of the crossing point less robust. Nevertheless, the conclusions drawn from the weak-coupling analysis carry over to the strong-coupling regime, as discussed in Section 5.2.1.

A series of reflection spectra is shown in Figure 6.9, that were obtained by modulating Δ_L while slowly varying Δ_p. Subfigures (a) and (e) represent the near-linear response by exhibiting only small perturbations of the primary-cavity resonance near degeneracy (marked by the intersection of the dashed white lines), indicative for weak coupling.

From these measurements, we extract a coupling strength of $g = 2\pi \times 68(5)$ MHz and a right-cavity linewidth of $\kappa_R = 2\pi \times 212(8)$ MHz.

From left to right, Figure 6.9 displays spectra measured with increasing input powers, spanning $P_{in} = 35(2)$ µW, 205(14) µW, 375(22) µW, and 627(34) µW. With increasing power the spectra develop progressively stronger asymmetries, characteristic of the thermally induced nonlinearity identified in the previous section.

Comment on the influence of the scan direction

The upper and lower rows of Figure 6.9 correspond to opposite scan directions of the length modulation. The shown spectra of opposite scan directions are qualitatively consistent at fixed input power, with only minor differences in the detailed lineshape. In the employed measurement scheme (detailed in Section 4.3.1), the system is probed by a continuously modulating ω_L while the probe frequency ω_p is approximately constant over a single cycle. Each sweep therefore starts from an initially unpopulated cavity, so that — within this protocol and parameter range — the response remains on the low-intensity branch of the auxiliary (right) cavity (see Figure 6.1 for the generic branch structure of a Kerr-type nonlinearity). The effective nonlinearities observed here are not sufficiently strong to trigger a transition to the high-intensity state; combined with the fact that the driven primary cavity remains approximately linear (and hence symmetric), the steady-state response is therefore expected not to exhibit optical hysteresis (bistability) under reversal of the scan direction. The residual scan-direction-dependent deviations are most plausibly attributed to dynamical effects associated with the non-uniform scan velocity, caused by the shifting the auxiliary-cavity resonance. This can distort the apparent cavity response during the scan (e.g., yielding an effective change in the linewidth inferred from the trace), in close analogy to Section 5.5, where we used a closely related mechanism to mimic an effective imbalance in the cavity linewidths. As a consistency check, we nevertheless retain both scan directions throughout the following analysis.

Nonlinear dispersion relation

From the two-dimensional spectra in Figure 6.9, we map the dispersion relation numerically by tracing the position of minimum reflectance as a function of ω_L for each ω_p. This corresponds to a horizontal slice-wise analysis of the two-dimensional reflection spectrum, analogous to the procedure used in Section 5.2. A minimum in the reflected signal corresponds to maximal coupling into the coupled system.

Figure 6.10 shows the extracted resonance condition for (a) positive and (b) negative scan rates, and for different input powers (represented by differently coloured traces). The solid black curve denotes the theoretical linear response expected for weak coupling to the auxiliary-cavity mode, for which the resonance condition (see Equation (5.4b),

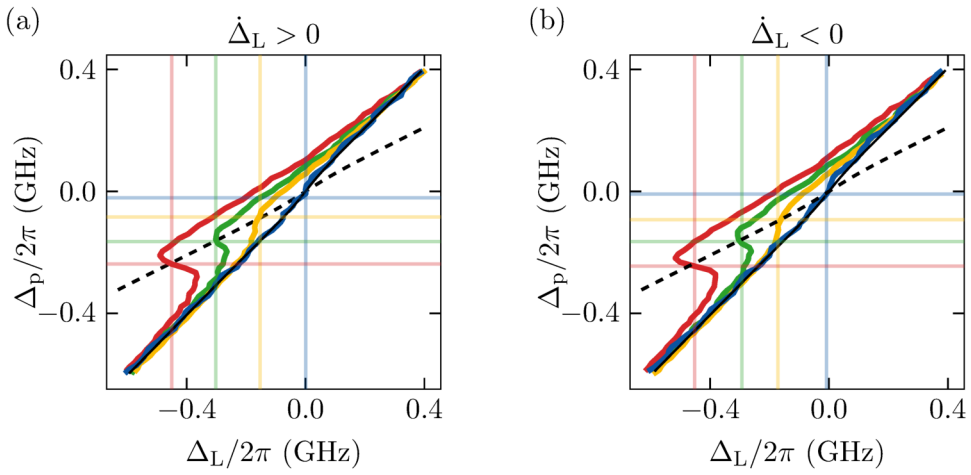


Figure 6.10.: (a) and (b) show the dispersion relation of the coupled system extracted from the resonances in Figure 6.9, plotted as a function of the left-cavity detuning Δ_L for both scan directions. Colours denote the incident powers $P_{in} = \{35(2)\ ,\ 205(14)\ ,\ 375(22)\ ,\ 627(34)\}$ µW. Solid black curves show the linear dispersion expected from the model, while dashed curves trace the inferred shift of the degeneracy point. The opaque horizontal markers indicate the measured maximal nonlinear shift of the right-cavity resonance, K'_{eff}. The corresponding vertical markers show the associated shift in Δ_L obtained by mapping this shift through the dispersion relation.

with $g = 2\pi \times 68(5)$ MHz and $\kappa_R = 2\pi \times 212(8)$ MHz) can be written as:

$$\omega_L - \omega_p = g^2 \frac{\omega_R - \omega_p}{(\omega_R - \omega_p)^2 + \left(\frac{\kappa_R}{2}\right)^2}. \tag{6.15}$$

The thermally induced Kerr-type response shifts the auxiliary-cavity resonance by an amount proportional to the intracavity photon number, which we parameterise as $K_{eff} \cdot |\tilde{\alpha}_R|^2$. Here, K_{eff} denotes the effective nonlinearity introduced in Section 6.1.1, encapsulating the power dependence of the resonance shift, while $|\tilde{\alpha}_R|^2$ quantifies the fraction of the incident field circulating in the auxiliary mode. Therefore, the nonlinear shift is maximal at the point of strongest coupling into the right cavity, which appears as a local minimum in the resonance amplitude along the extracted dispersion trace (see Equation (5.4c)). This point can be identified and related to the probe detuning Δ_p via:

$$\Delta_p = \omega_p - \omega_R = \underbrace{K_{eff} \cdot \underbrace{\max\left(|\tilde{\alpha}_R|^2\right)}_{\frac{4g^2\kappa_L^2}{(4g^2+\kappa_L\kappa_R)^2}}}_{K'_{eff}}. \tag{6.16}$$

The intersections of the coloured horizontal and vertical lines indicated in Figure 6.10 mark the corresponding coordinates (Δ_L, Δ_p) of maximal coupling to the auxiliary-cavity mode and thus provide an experimental estimate of K'_{eff} for each input power.

To relate the observed vertical shift in the two-dimensional spectra (a change in Δ_p) to the corresponding horizontal shift of the coupled resonance (a change in Δ_L), we evaluate the linear dispersion relation at $\Delta_\mathrm{p} = K'_\mathrm{eff}$ while accounting for the auxiliary-cavity resonance shift of the auxiliary cavity ($\omega_\mathrm{R} \to \omega_\mathrm{R} + K'_\mathrm{eff}$). This yields an effective shift of the degeneracy point along the horizontal axis, given by:

$$\begin{aligned} \Delta_\mathrm{L} - K'_\mathrm{eff} &= \left(1 - \frac{g^2}{(K'_\mathrm{eff})^2 + (\kappa_\mathrm{R}/2)^2}\right) K'_\mathrm{eff} \\ \iff \qquad \Delta_\mathrm{L} &= \left(2 - \frac{g^2}{(K'_\mathrm{eff})^2 + (\kappa_\mathrm{R}/2)^2}\right) K'_\mathrm{eff}, \end{aligned} \tag{6.17}$$

where Δ_L remains defined relative to the unshifted auxiliary-cavity resonance ω_R. The dashed curves in Figure 6.10 show Δ_L as a function of K'_eff and agree well with the extracted intersection points, indicating that the measured drift of the degeneracy point is consistent with the above mapping.

Because $|\tilde{\alpha}_\mathrm{R}|^2$ depends nonlinearly on Δ_L and Δ_p, reconstructing the full nonlinear dispersion beyond the degeneracy point generally requires numerical modelling to capture the self-consistent interplay between intracavity field amplitude and detuning. Nevertheless, the shift of the degeneracy point already captures the key back-action mechanism of interest here: the auxiliary nonlinearity maps onto an effective detuning of the coupled resonance through the system's dispersion relation.

This dispersive back-action can equivalently be expressed as a photon-number-dependent phase shift imprinted by the auxiliary right cavity onto the driven primary left-cavity response. In the following, we use this phase picture to derive the optimal parameter regime in which two-photon excitation pathways acquire opposite phases and interfere destructively, thereby resulting in unconventional photon blockade.

6.2.2. Optimal parameters for destructive interference

Within the semi-classical description sketched in Figure 6.8.(b), the interference condition can be formulated as a requirement on the effective detuning of the primary-cavity response for different auxiliary-cavity states. Specifically, destructive interference requires that the effective primary left-cavity resonance is detuned by the same amount, but in opposite directions, depending on whether the auxiliary right cavity is unshifted (empty) or shifted by the single-photon nonlinearity K (populated). Using the standard complex reflection coefficient of a single cavity, detuning of $\omega_\mathrm{L} - \omega_\mathrm{p} = \pm\kappa_\mathrm{L}/2$ correspond to opposite phases of the reflected field amplitude (and hence opposite signs of the intracavity phase), with a relative phase shift of π (up to sign convention) [30].

For an unshifted auxiliary cavity, the coupled-resonance condition therefore requires:

$$\omega_\mathrm{L} - \omega_\mathrm{p} = g^2 \frac{\omega_\mathrm{R} - \omega_\mathrm{p}}{(\omega_\mathrm{R} - \omega_\mathrm{p})^2 + \left(\frac{\kappa_\mathrm{R}}{2}\right)^2} \overset{!}{=} \pm\frac{\kappa_\mathrm{L}}{2}, \tag{6.18}$$

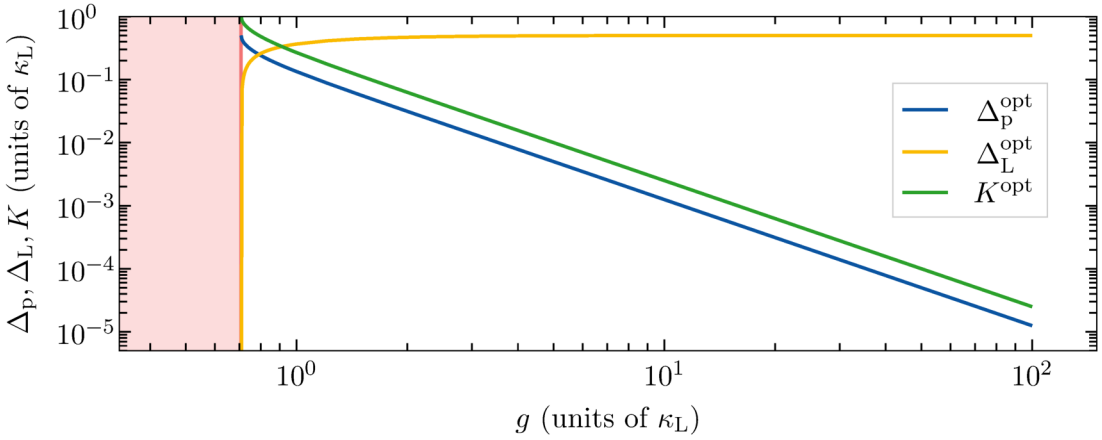


Figure 6.11.: Optimal parameters for the probe detuning $\Delta_\mathrm{p} = \omega_\mathrm{p} - \omega_\mathrm{R}$, cavity–cavity detuning $\Delta_\mathrm{L} = \omega_\mathrm{L} - \omega_\mathrm{R}$, and single-photon nonlinearity K as functions of the coupling strength g, obtained from the destructive-interference conditions. These parameters enable suppression of two-photon excitation in the driven (left) cavity via interference, i.e. unconventional photon blockade. The red-shaded region indicates the parameter range where no real-valued solutions exist and the interference condition cannot be satisfied.

whereas a population-dependent shift of the auxiliary mode, $\omega_\mathrm{R} \to \omega_\mathrm{R} + K$, leads to:

$$\omega_\mathrm{L} - \omega_\mathrm{p} = g^2 \frac{\omega_\mathrm{R} - \omega_\mathrm{p} + K}{(\omega_\mathrm{R} - \omega_\mathrm{p} + K)^2 + (\frac{\kappa_\mathrm{R}}{2})^2} \stackrel{!}{=} \mp \frac{\kappa_\mathrm{L}}{2}. \tag{6.19}$$

Under these two conditions, the circulating field acquires opposite roundtrip phases in the two cases, fulfilling the destructive-interference criterion underlying the unconventional photon blockade.

Solving the above equations yields optimal values[5] for the probe detuning $\Delta_\mathrm{p} = \omega_\mathrm{p} - \omega_\mathrm{R}$, the cavity–cavity detuning $\Delta_\mathrm{L} = \omega_\mathrm{L} - \omega_\mathrm{R}$, and the single-photon nonlinearity K as functions of the coupling strength g:

$$\Delta_\mathrm{p}^\mathrm{opt} = \left(\frac{g^2}{\kappa_\mathrm{L}}\right) - \sqrt{\left(\frac{g^2}{\kappa_\mathrm{L}}\right)^2 - \left(\frac{\kappa_\mathrm{R}}{2}\right)^2}, \tag{6.20a}$$

$$\Delta_\mathrm{L}^\mathrm{opt} = \frac{\kappa_\mathrm{L}}{2} - \Delta_\mathrm{p}^\mathrm{opt}, \tag{6.20b}$$

$$K^\mathrm{opt} = 2\Delta_\mathrm{p}^\mathrm{opt}. \tag{6.20c}$$

These dependencies are illustrated in Figure 6.11, normalised to κ_L and assuming $\kappa_\mathrm{L} = \kappa_\mathrm{R}$. Real-valued solutions exist only if the discriminant is positive, i.e. for $g \geq$

[5]Here, we focus on the positive solutions, despite the photothermal shift corresponds to an effective $K < 0$. This choice facilitates a direct comparison of the resulting scaling behaviour with the quantum master-equation treatment in Ref. [37]. However, the parameters are invariant under redefinition of detuning/phase.

$\sqrt{(\kappa_L \kappa_R)/2}$, indicated by the red-shaded unavailable region. In the strong-coupling limit, the required nonlinearity decreases as:

$$K^{\text{opt}} = \frac{1}{2}\frac{\kappa_L \kappa_R^2}{g^2} + \mathcal{O}(g^{-4}), \tag{6.21}$$

indicating that the interference condition can, in principle, be satisfied even for small single-photon nonlinearities at sufficiently large g. This contrasts with conventional photon blockade in a single nonlinear cavity, where suppressing two-photon excitation spectrally typically requires an interaction-induced resonance shift comparable to (or larger than) the cavity linewidth [37].

Comparison to quantum-theoretical results

In Ref. [37], the unconventional photon blockade in a coupled-cavity system is analysed by solving the quantum master equation, including excitation processes up to the two-photon manifold. Overall, the optimal parameter trends obtained from this full quantum treatment are consistent with the semi-classical criteria derived in Section 6.2.2, but two differences are particularly relevant.

First, close to the boundary of the inaccessible region at small coupling strengths (red-shaded area in Figure 6.11), the semi-classical model predicts an approximately constant optimal nonlinearity. In the master-equation solutions, by contrast, the required nonlinearity increases sharply in this regime. A natural interpretation is that when the single-photon pumping rate (here taken as $F = \kappa_L/2$) becomes comparable to the inter-cavity coupling rate g, the two excitation pathways sketched in Figure 6.8 no longer evolve on similar timescales. This temporal mismatch reduces the degree of coherent interference between the pathways and therefore reduces the destructive-interference contrast, again requiring substantially larger nonlinearities. The semi-classical model, which effectively assumes instantaneous coherent exchange compared to the excitation dynamics, does not capture this effect and consequently underestimates the required nonlinearity near the boundary. In the strong-coupling limit $g \gg \kappa_L/2$, both approaches converge and yield the same qualitative behaviour, in particular $K^{\text{opt}} \propto g^{-2}$ with an approximately constant optimal detuning.

Second, the quantum treatment accounts for additional contributions arising from intermediate occupation of higher-lying states, most notably $|02\rangle$ (light blue in Figure 6.8.(a)). These processes introduce an additional phase contribution to the two-photon excitation amplitude and lead to a quantitative correction of the optimal nonlinearity:

$$K^{\text{opt}} \approx \frac{2}{3\sqrt{3}}\frac{\kappa_L \kappa_R^2}{g^2}, \tag{6.22}$$

while preserving the same g^{-2} scaling.

Despite these simplifications, the semi-classical framework captures the physical origin of unconventional photon blockade — namely interference between two-photon excita-

tion pathways — and provides an intuitive route to estimating favourable parameter regimes. The agreement in scaling with the quantum treatment supports the viability of implementing the effect in the coupled-cavity platform considered here, subject to the dynamical constraints discussed below.

6.2.3. Prospects for experimental implementation

A key practical constraint for unconventional photon blockade is that the nonlinearity must act on the timescale of the cavity dynamics. In the present implementation, the dominant nonlinearity of the silicon nitride layer is photothermal, with a relaxation time $\tau_{\mathrm{T}} \approx 25$ µs (Section 6.1.2), i.e. many orders of magnitude slower than both κ_{L}^{-1} and g^{-1}. As a result, the auxiliary-cavity resonance cannot acquire a photon-number–dependent shift on the timescale on which photons enter and leave the coupled system, and a second photon may couple in before the resonance has responded, thereby penetrating the blockade.

By contrast, the electronic Kerr effect is effectively instantaneous on optical timescales and would, in principle, satisfy the dynamical requirements [130]. Using Equation (6.22) (including the quantum-mechanical correction), the single-photon nonlinearity required for destructive interference is:

$$K^{\mathrm{opt}} \approx 2\pi \times 0.3\ \mathrm{MHz}, \tag{6.23}$$

for $g = 2\pi \times 2.86$ GHz, $\kappa_{\mathrm{L}} = 2\pi \times 212$ MHz, and $\kappa_{\mathrm{R}} = 2\pi \times 168$ MHz, as determined in Section 4.3. To invoke a detuning of this strength with only a single photon, we require an effective nonlinear constant β of:

$$\beta = \frac{K^{\mathrm{opt}}}{P_{\mathrm{in}}} \approx K^{\mathrm{opt}} \cdot \frac{c}{L_{\mathrm{R}} \cdot \hbar\omega} = 2\pi \times 3.49\ \frac{\mathrm{GHz}}{\mathrm{mW}}, \tag{6.24}$$

which is far above the photothermal β derived in Section 6.1.2. With Equations (6.8) and (6.12), this translates to a required nonlinear refractive index of approximately:

$$n_2 \approx 1.1 \times 10^{-14}\ \frac{\mathrm{m}^2}{\mathrm{W}}. \tag{6.25}$$

well beyond the Kerr nonlinearities of conventional dielectrics [120–122]. In particular, silicon nitride remains insufficient in the present geometry, in part because the film contributes only weakly to the overall mode volume ($t_{\mathrm{SiN}} \approx 0.5$ µm versus $d_{\mathrm{R}} \approx 150$ µm).

A natural route to increase the effective nonlinearity is therefore to fill the cavity with a nonlinear medium. Liquid-filled cavities with high optical quality have recently been demonstrated in our group [29]; in that work, a synthetic silicone oil with $n_2 = 3.62(31) \times 10^{-20}$ m^2/W occupied the full cavity volume. While this specific liquid may remain below the value required here, it establishes the feasibility of high-finesse liquid-filled resonators. More broadly, the material space of nonlinear liquids (and mixtures) is large, offering media that combine a stronger Kerr response with sufficient response

times and low absorption [131].

However, if the nonlinearity is distributed throughout the cavity volume, the required index nonlinearity is reduced to:

$$n_2 \approx 7 \times 10^{-17}\ \mathrm{m}^2/\mathrm{W}. \tag{6.26}$$

which is still about three orders of magnitude above the value reported in [29]. Importantly, this requirement is highly sensitive to cavity linewidth: in the present scaling, $K^{\mathrm{opt}} \propto \kappa_{\mathrm{R}}^2$ while the intracavity enhancement scales with finesse $\mathcal{F}_{\mathrm{R}} \propto \kappa_{\mathrm{R}}^{-1}$, resulting in an overall dependence that is approximately cubic in κ_{R}. Thus, improvements in κ_{R} can substantially relax the required material nonlinearity.

In this context, it is important to note that an earlier state of the apparatus (Section 4.2.1) achieved a significantly higher finesse for the left cavity ($\kappa_{\mathrm{L}} = 2\pi \times 28$ MHz) before the linewidth later degraded. Crucially, in a liquid-filled implementation the nonlinearity would be provided by the intracavity medium rather than by the silicon nitride membrane, decoupling the nonlinear response from the present assignment of the right cavity as the nonlinear site and making it natural to interchange the roles of primary and auxiliary cavities based on their optical performance. If the higher-finesse configuration could be recovered and the cavity roles are swapped, the required n_2 decreases further (by approximately a factor ~ 200 in our estimate), reducing the remaining gap to using the tested synthetic silicone oil to roughly an order of magnitude.

Finally, the inter-cavity coupling g provides an additional lever. Increasing the membrane transmittance is an obvious route; although this would increase intracavity losses in a naive single-cavity picture. Figure 4.6 indicates that the effective reflection of the secondary cavity can exceed that of the bare membrane mirror and can be limited primarily by the fibre mirror transmittance, suggesting that κ_{L} and κ_{R} do not change proportionally. However, at the same time, strong asymmetries in end-mirror reflectivity increase impedance mismatch and can reduce coupling efficiency offsetting potential gains. As an alternative (or complement), increasing the cavity length can extend the interaction length while keeping the ratio g/κ_{L} approximately constant, thereby relaxing the required nonlinearity through increased effective interaction volume.

Taken together, these considerations point to a realistic path toward an experimental demonstration of the unconventional photon blockade in a future iteration of the setup: by combining a faster, Kerr-dominated nonlinear medium (potentially in a liquid-filled implementation) with improved linewidths/finesse and an increased inter-cavity coupling g, the destructive-interference condition identified above could plausibly be brought within reach employing components that are already available and validated in the present platform.

6.3. Summary

Absorption-induced heating in the silicon nitride layer of the membrane mirror introduces a photothermal nonlinearity into the coupled-cavity platform: a fraction of the

circulating field is absorbed, raising the local temperature and thereby changing the optical thickness of the membrane. This produces an intensity-dependent detuning and extends the photonic-molecule picture developed in the previous chapter into a regime where the coupled resonances exhibit measurable nonlinear back-action.

To characterise the origin and strength of this effect, we first studied the right cavity containing the silicon nitride layer while operating the left cavity off resonance. In this configuration, the nonlinearity manifests as optical hysteresis and bistability, such that identical drive conditions can yield different cavity states depending on the system history. By mapping the bistable region as a function of incident power, we quantified the effective photothermal tuning strength as $\beta = -2\pi \times 13.8(1)\,\mathrm{MHz/mW}$.

We then investigated the temporal response of the nonlinearity to further test the thermal interpretation. By measuring the bistable window for different scan rates and powers, we inferred a characteristic response time on the order of $\tau_\mathrm{T} \sim 25\,\mu\mathrm{s}$. The slow, scan-rate-dependent dynamics strongly support a photothermal mechanism as the dominant origin of the observed nonlinear behaviour.

In the coupled-cavity configuration, we connected this nonlinear response to the power-dependent dispersion of the coupled system extracted from two-dimensional reflection spectra. In the weak-coupling regime ($g = 2\pi \times 68(5)$ MHz, $\kappa_\mathrm{L} = 2\pi \times 212(2)$ MHz, $\kappa_\mathrm{R} = 2\pi \times 212(8)$ MHz), this approach provides a direct view of how the nonlinearity of the right cavity shifts the degeneracy point and imprints phase back-action on the effective resonance of the coupled system. This dispersion-based viewpoint naturally motivates an interference picture for photon-blockade-like behaviour: unlike conventional (spectral) photon blockade, which requires a nonlinear shift exceeding the linewidth ($K \gg \kappa_\mathrm{R}$), the unconventional photon blockade (UPB) suppresses two-photon excitation primarily through destructive interference between distinct excitation pathways.

Using the experimentally extracted dispersion, we developed a simplified semi-classical model that predicts the optimal parameters for UPB in the strong-coupling limit and compared the resulting scaling to a full quantum master-equation treatment [37]. For the strongly coupled mode pair used here ($g = 2\pi \times 2.86$ GHz, $\kappa_\mathrm{L} = 2\pi \times 212$ MHz, $\kappa_\mathrm{R} = 2\pi \times 168$ MHz), this yields an estimate of the required single-photon nonlinearity of $K^\mathrm{opt} \approx 2\pi \times 0.3$ MHz, i.e. far below the cavity linewidth.

Despite this substantial reduction compared to the conventional blockade condition, achieving a Kerr-dominated nonlinearity of the required magnitude remains challenging for standard materials. Moreover, the interference-based approach requires coherence between excitation pathways on the timescale of the cavity dynamics, whereas the photothermal response of the silicon nitride membrane is orders of magnitude slower ($\tau_\mathrm{T} \gg \kappa_\mathrm{L}^{-1}, g^{-1}$). We therefore outlined practical routes toward a future demonstration, including liquid-filled cavities [29] and improvements of linewidth and inter-cavity coupling that can be pursued using the same class of components and assembly approach as employed throughout this thesis. Taken together, these results suggest that the coupled-cavity architecture developed here provides a realistic and scalable basis for approaching an experimental observation of unconventional photon blockade in a future iteration of the setup.

7. Summary and Future Direction

This thesis presents the experimental implementation and characterisation of two coherently coupled, fibre-integrated Fabry–Pérot resonators as a versatile platform for studying dissipative quantum optics. By enabling photon tunnelling between two optical microcavities, the system realises a photonic molecule whose hybridised modes provide the natural basis for interference and loss-engineering phenomena in coupled-resonator settings [15, 16]. The device is based on a single Fabry–Pérot cavity formed between the end facets of two opposing single-mode fibre mirrors [9, 51]. Introducing a high-reflectivity membrane mirror into this cavity divides the original resonator into two cavities of equal length (150 µm), thereby creating two coupled cavity sites. The membrane mirror controls coherent coupling between the cavities via photon transmission.

A key contribution of this work was the development of the membrane mirror, fabricated on suspended silicon nitride membranes with an optimised dielectric coating. Stress management through the integration of a high-stress layer, together with a 200 µm aperture mask, mitigated deformation of the mirror surface after coating deposition. As a result, the achieved surface quality and optical performance of the membrane mirror (radius of curvature $R \approx 8$ mm, transmittance $\mathcal{T} = 452(47)$ ppm) were sufficient to reliably sustain high-finesse cavities on both sides.

The fundamental TEM_{00} modes, identified by efficient coupling to the fibre waveguides, exhibit linewidths of $\kappa_\mathrm{L} = 2\pi \times 212(2)$ MHz and $\kappa_\mathrm{R} = 2\pi \times 168(3)$ MHz, respectively. High-resolution spectroscopy near the degeneracy point resolved an avoided crossing with coupling strength $g = 2\pi \times 2.859(2)$ GHz (defined as half the normal-mode splitting on resonance). The measured coupling strength is 20 % below the theoretical prediction, which we attributed to slight cavity misalignment. This imperfection also lifts the orthonormality of transverse cavity modes of different order, which we exploited to tune the coupling strength over nearly three orders of magnitude, spanning the weak- to strong-coupling regimes, beyond the limit implied by the transmittance of the membrane mirror.

Within this all-optical setting, we observed clear light–matter analogues — reinforcing the photonic-molecule picture — including Purcell-like rate modification, normal-mode splitting, and EIT-like transparency [14, 20, 100]. Since true EIT typically requires a decay-rate asymmetry [101], we emulated the necessary imbalance by dynamically modulating the individual cavity lengths. This enabled transparency features in a purely optical, lossy two-mode system without requiring a metastable state.

We further identified an intrinsic nonlinearity in one cavity arising from absorption in the silicon-nitride membrane. This nonlinearity produces optical bistability and hysteresis [59], which we characterised experimentally. In the coupled system, the nonlinearity

extends the photonic molecule into the nonlinear regime, such that transparency effects become photon-number dependent. Based on these measurements and analysis, we discussed realistic conditions under which destructive interference can be engineered to suppress the two-photon excitation pathway, enabling the unconventional photon blockade mechanism [34, 37, 117] and providing a concrete outlook for future experiments.

Overall, we demonstrated that this platform provides direct experimental control over coherent coupling, (emulated) decay imbalance, and material nonlinearity, opening a broad parameter space for dissipation engineering and nonlinear photonics in the steady-state limit. Furthermore, transverse access to one cavity site enables the integration of additional nonlinear elements (e.g., single atoms or nonlinear liquids), connecting this work to prospects in quantum simulation with interacting photons [32, 33].

7.1. Outlook: Temporally resolved photon tunnelling between cavity sites

So far, our analysis has been largely restricted to steady-state observables. Even in the discussion of the nonlinear photothermal response time in Section 6.1.3, the relevant dynamics evolve slowly compared with the rapid cavity-field decay on the timescale $\mathcal{O}(\kappa_{\mathrm{L/R}}^{-1})$. By contrast, coherent photon exchange between the two cavities occurs on the coupling timescale set by g, which is substantially faster.

In the time domain, this coherent exchange leads to oscillatory energy transfer (photon tunnelling) superimposed on the exponential ringdown of the intracavity energy. However, if the system is first loaded into its (quasi-)steady state by continuous-wave driving and the drive is then switched off, both sub-cavities start from nearly equal populations. In this case, the bidirectional coherent exchange produces little net population modulation in either cavity, so the oscillatory tunnelling signal exhibits negligible contrast. Therefore, to reveal coherent tunnelling, one must prepare an initial population imbalance. Two natural routes are possible: (1) interference-based state preparation that selectively suppresses one hybrid mode via controlled detuning and linewidth balance [95], and (2) impulsive excitation of a single sub-cavity using a short optical pulse of duration $\delta t \lesssim g^{-1} \sim 50$ ps.

Figure 7.1 compares the resulting decays of the cavity populations for (a) an equilibrium initial state and (b) a state prepared by a short optical pulse. In both cases the population follows the same decay envelope, while the oscillation contrast depends sensitively on the prepared initial condition. For this simulation, we solve the equation of motion, see Equation (4.4), numerically for a temporally varying drive field $f(t)$:

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix}\alpha_{\mathrm{L}}\\ \alpha_{\mathrm{R}}\end{pmatrix} = \begin{pmatrix} i(\Delta_{\mathrm{L}}-\Delta_{\mathrm{p}})-\frac{\kappa_{\mathrm{L}}}{2} & -ig \\ -ig & i\Delta_{\mathrm{p}}-\frac{\kappa_{\mathrm{R}}}{2}\end{pmatrix}\cdot\begin{pmatrix}\alpha_{\mathrm{L}}\\ \alpha_{\mathrm{R}}\end{pmatrix} + f(t)\begin{pmatrix}1\\ 0\end{pmatrix}. \tag{7.1}$$

From an implementation standpoint, the interference route requires precise combinations of κ_{L}, κ_{R}, g, and detuning, which restricts the accessible parameter space. The impulsive route is experimentally more direct, provided that sufficiently short pulses are available.

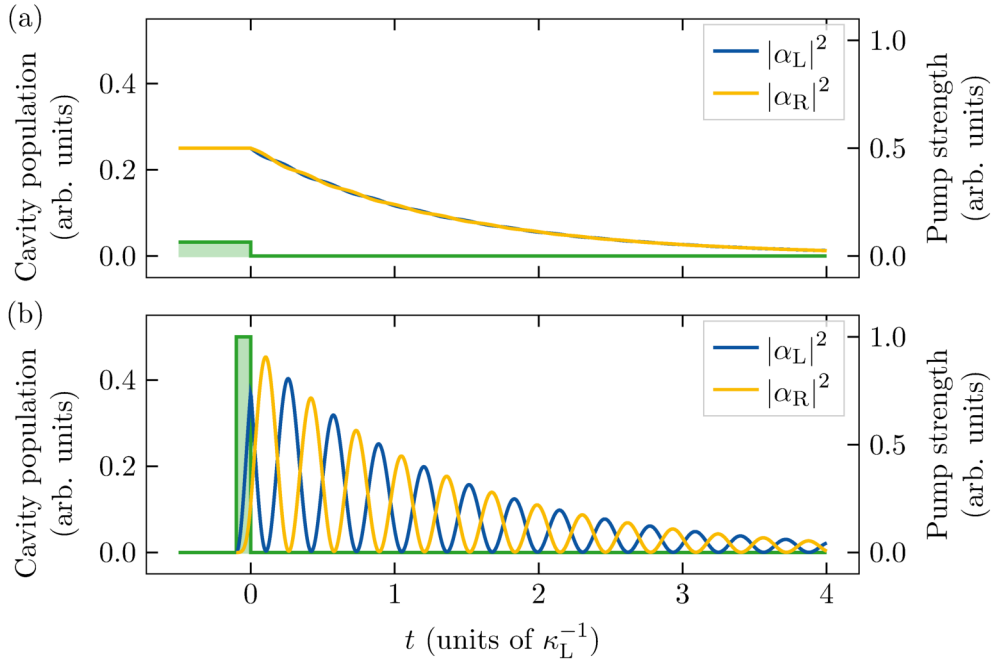


Figure 7.1.: Simulation of the ringdown of the intracavity population, with the initial state being prepared (a) by continuous pumping and (b) with an ultra-short rectangular optical pulse with a width of δt. The (relative) pump strength (green) is shown on the right axis. The superimposed modulation is caused by the coherent exchange of population between both cavity sites. The simulation assumes the following parameters: $(\kappa_R/\kappa_L) = 0.5$, $(g/\kappa_L) = 10$, $\Delta_L = 0$, $\Delta_p = g$, and $\delta t = g^{-1}$.

The electro-optic amplitude modulator introduced in Section 6.1.2 (specified rise time < 5 ps), together with a fast electrical driver (built and tested by Dr. Akos Hoffmann[1]), offers a realistic path to optical pulses with δt well below the inter-cavity exchange timescale. In combination with fast detection on both sides of the cavity, an ultra-short pump pulse should make time-resolved decay of the exponential ring-down of the individual cavities (directly reflecting the respective cavity populations) experimentally accessible with the current setup. The revealed superimposed modulation then forms a photonic analogue of Rabi oscillations [16]. With reduced input power and improved detection efficiency, operation at or near the single-photon level would bring the real-time observation of photon tunnelling between the two cavities within reach. Such a time-resolved measurement would directly verify the exchange rate inferred from steady-state spectra and provide controlled access to coherent dynamics in a deliberately lossy photonic environment [132].

Beyond benchmarking coherent exchange, access to the few-photon dynamics would also sharpen the experimental signature of interference-based nonlinear effects. In particular, it would provide direct information on the temporal build-up of the unconventional

[1]Contact: ahoffman@physik.uni-bonn.de

photon blockade discussed in Section 6.2: because this effect relies on the timed cancellation of distinct two-photon excitation pathways, its unambiguous identification requires resolving the photon population on the inter-cavity exchange timescale (g^{-1}), rather than only in the stationary limit [36, 37, 133].

Taken together, these considerations outline a clear experimental path: extending the platform from steady-state spectroscopy to time-domain measurements enables a direct probe of coherent tunnelling under controlled dissipation, while simultaneously providing the temporal resolution to characterise the interference-based few-photon nonlinearities. Combined with the promising route outlined in the previous chapter, where we considered filling one cavity site with a nonlinear liquid to enhance both the effective nonlinearity and the relevant response time, this outlook already contains a realistic prospect for achieving unconventional photon blockade in a future iteration of the setup.

Overall, time-resolved access to inter-cavity exchange complements these nonlinear upgrade paths by opening a route to extend interacting-photon quantum simulation from steady-state observables to engineered non-equilibrium dynamics in a controlled, driven dissipative two-mode setting [32, 33, 39, 134]. Moreover, nonclassical light, as can be generated via the unconventional photon blockade [35], is a key resource for quantum communication and photonic information processing. These results therefore motivate an interference-based avenue toward photonic quantum technologies, particularly in settings where nonclassical light can be generated and controlled within compact resonator platforms [38].

A. Introduction to the transfer-matrix method

Contents

A.1. Reflection and transmission spectra of the multilayer stack . 119
A.2. Intensity distribution within the mirror coating 123

In Chapter 2, we derived the reflection and transmission spectra of a cavity resonator by representing the forward-propagating (incident or transmitted) and backward-propagating (reflected) field components of an optical wave as a two-component field vector $\boldsymbol{a}_i$. Within this framework, interference effects were analysed by decomposing the cavity into its constituent optical elements, each of which maps an incident field vector to an output field vector and thus propagates the field through the system.

A natural extension of this idea is the transfer-matrix formalism, which provides a robust and convenient framework for tracking the evolution of the field vector through complex, multilayer structures. In particular, it enables efficient prediction of reflection, transmission, and internal field amplitudes for layered optical systems. This makes it especially useful for modelling dielectric coating stacks, such as the multilayer coating on the silicon nitride membrane introduced in Chapter 3.

In this section, we build on the formalism developed in Chapter 2 and introduce the generalised transfer-matrix approach. We then demonstrate basic simulations for simple dielectric multilayers composed of $\lambda/4$ layers. Finally, we outline how the method can be extended from global spectra to a spatially resolved description of the field inside the stack, allowing the intensity distribution within each layer to be computed. This capability is central to the subsequent analysis of absorption in the membrane mirrors (see Section 3.2.2).

The presentation follows Ref. [41], with notation adapted for consistency with the remainder of this thesis.

A.1. Reflection and transmission spectra of the multilayer stack

A multilayer stack consists of a sequence of dielectric layers with different refractive indices. When an optical field interacts with such a structure, partial reflections arise at every interface where the refractive index changes. These reflections follow from the electromagnetic boundary conditions and are quantified by the Fresnel equations [43]. For oblique incidence, the field components parallel and perpendicular to the plane of

incidence must be treated separately. At normal incidence, however, this distinction is not required, and the complex amplitude reflection and transmission coefficients, ρ and τ, reduce to:

$$\rho = \frac{n_i - n_f}{n_i + n_f}, \tag{A.1a}$$

$$\tau = \frac{2n_i}{n_i + n_f}, \tag{A.1b}$$

where n_i and n_f denote the refractive indices of the initial and final media, respectively. Derivations for arbitrary angles of incidence can be found in the literature (e.g. Ref. [41] or Ref. [43]) and are not repeated here.

If the partial reflections from successive interfaces add coherently, constructive interference can strongly enhance the overall reflectance. Conversely, destructive interference suppresses reflection and increases transmission. The optical thickness of each layer therefore plays a central role, as it determines the interference condition at the wavelength of interest. This is closely analogous to a Fabry–Perot resonator, whose spectral response is governed by the round-trip phase accumulated in the cavity (see Section 2.1). From this perspective, each layer may be viewed as a low-finesse resonator, and the entire coating as a cascade of such resonant sections.

The transfer matrix that describes the propagation of a field vector across the boundary between the n-th and $(n+1)$-th layer is given by:

$$\mathcal{M}_n = \frac{1}{\tau_n} \begin{pmatrix} 1 & -\rho_n \\ \rho_n & 1 \end{pmatrix}, \tag{A.2a}$$

where the complex Fresnel coefficients ρ_n and τ_n are defined as in Equation (A.1). Compared to the interface matrices introduced in Section 2.1, the relative sign between the off-diagonal terms differs, reflecting the sign convention implied by the Fresnel relations when reversing the direction of incidence (i.e. considering reflection from opposite sides). The normalisation $\mathcal{M}_{11} = 1$ ensures that the matrix represents a lossless boundary transformation (i.e. no light is lost at the interface).

Propagation through the n-th layer between adjacent interfaces is captured by the diagonal phase matrix:

$$\mathcal{P}_n = \begin{pmatrix} \mathrm{e}^{-i\varphi_n} & 0 \\ 0 & \mathrm{e}^{i\varphi_n} \end{pmatrix}, \tag{A.2b}$$

where φ_n is the phase accumulated over a single pass through the layer: $\varphi_n = n_n t_n k_0$. Here, n_n and t_n are the refractive index and thickness of the n-th layer, respectively, and $k_0 = 2\pi/\lambda_0$ is the vacuum wavenumber. In general, n_n may be complex, where an imaginary component accounts for attenuation during propagation and thus provides a convenient model of absorption within the layer.

The total transfer matrix of an N-layer stack is obtained by multiplying the interface

and propagation matrices in sequence. The product is ordered along the optical path through the stack: starting from the incident medium, each subsequent interface and propagation step is appended, so that the overall matrix acts on the field vector from right to left. The resulting system matrix $\mathcal{S}$ can be written as:

$$\begin{aligned}\mathcal{S} &= \mathcal{M}_N \otimes \mathcal{P}_{N-1} \otimes \mathcal{M}_{N-1} \otimes \ldots \otimes \mathcal{M}_0 \otimes \mathcal{P}_0 \otimes \mathcal{M}_{-1} \\ &= \left(\prod_{i=N}^{0} \mathcal{M}_i \otimes \mathcal{P}_i \right) \otimes \mathcal{M}_{-1},\end{aligned} \tag{A.3}$$

where $\mathcal{M}_{-1}$ describes the transition from the ambient medium into the first layer and $\mathcal{M}_N$ the transition from the final layer into the exit medium. Analytical expressions for $\mathcal{S}$ can be derived for small stacks, but the algebra quickly becomes complicated as the number of layers increases. In practice, $\mathcal{S}$ is therefore evaluated numerically, which is efficient even for large stacks.

To compute the reflectance and transmittance of the multilayer system, we note that $\mathcal{S}$ represents a mirror-like optical element that partitions incident light into reflected and transmitted components, and — if any layer has a complex refractive index ($n_k \in \mathbb{C}$) — an absorbed component. Using the same conventions as in Equation (A.2a), the power reflectance and transmittance are:

$$\mathcal{R}_\mathcal{S} = \left| \frac{\mathcal{S}_{21}}{\mathcal{S}_{22}} \right|^2, \tag{A.4a}$$

$$\mathcal{T}_\mathcal{S} = \frac{n_{N+1}}{n_{-1}} \times \left| \frac{1}{\mathcal{S}_{22}} \right|^2. \tag{A.4b}$$

where the prefactor n_{N+1}/n_{-1} accounts for differences in energy density between the incident and exit media. This factor is particularly relevant for coatings deposited on a substrate, as in Bragg mirrors or anti-reflection (AR) stacks, where typically $n_{N+1} \neq n_{-1}$. In the present work, the silicon nitride layer is embedded within a film that is effectively surrounded by air (or vacuum) on both sides, such that the ambient medium is the same on both sides of the stack.

Dielectric multilayer stacks serving as Bragg mirror

A Bragg mirror is a high-reflectivity dielectric mirror formed by alternating layers of low and high refractive index. Its large reflectance arises because the partial fields reflected at successive interfaces can be engineered to recombine constructively at a design wavelength λ_0 [43].

For constructive interference, the reflected contributions must add with a net relative phase equal to an integer multiple of 2π. In an alternating stack, successive interface reflections originate from opposite index steps, such that every other partial reflection occurs at the optically denser side of an interface and therefore acquires an additional phase shift of π (i.e., $\rho_n = -\rho_{n+1}$). Consequently, the propagation phase accumulated between successive reflection events must supply the remaining phase, implying that the

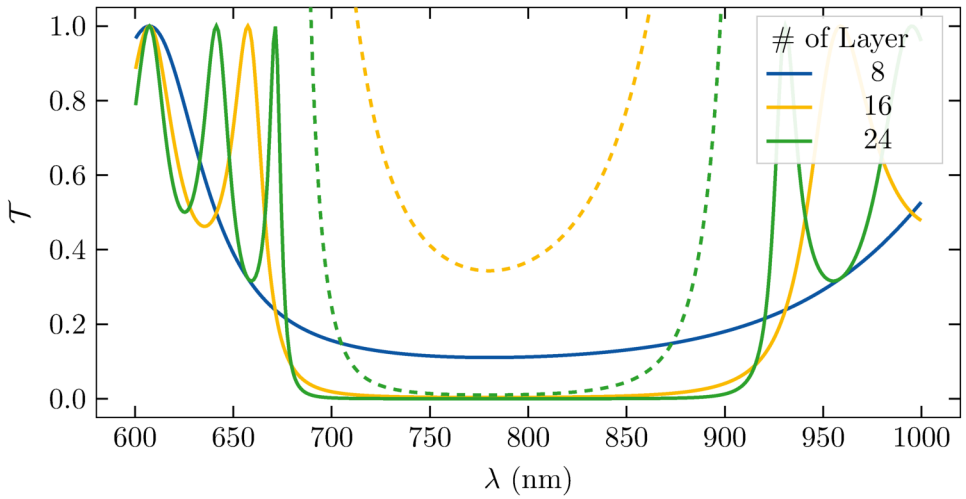


Figure A.1.: Transmittance $\mathcal{T}$ of SiO_2/Nb_2O_5 quarter-wave Bragg stacks with $N = \{8, 16, 24\}$, designed for $\lambda_0 = 780$ nm. Dashed curves are scaled by 100.

round-trip propagation phase in each layer is an odd multiple of π. This is achieved when the layer thickness t_n corresponds to an odd integer multiple q of the quarter-wavelength optical thickness, i.e.

$$t_n = q\left(\frac{\lambda_n}{4}\right) = q\left(\frac{\lambda_0}{4n_n}\right) \equiv q \text{ QWOT}. \tag{A.5}$$

Because of its central role in Bragg mirror design, layer thickness is commonly specified in units of the quarter-wavelength optical thickness (QWOT).

For a multilayer stack composed of N alternating layers with refractive indices n_1 and n_2 (with $n_2 > n_1$), and with each layer having an optical thickness of 1; QWOT, the reflectance can be approximated as:

$$\mathcal{R}_{\text{QWOT}} \approx 1 - 4\left(\frac{n_1}{n_2}\right)^N \begin{cases} \frac{1}{n_f} & \text{for } N \text{ even} \\ \frac{n_f}{n_1 n_2} & \text{for } N \text{ odd} \end{cases} \tag{A.6}$$

where n_f denotes the refractive index of the final medium (typically the substrate) [42]. In this configuration, the reflectance approaches unity with a comparatively small number of layers.

As an example, consider a Bragg mirror consisting of alternating layers of silicon dioxide (SiO_2) and niobium pentoxide (Nb_2O_5), with refractive indices $n_{\text{SiO2}} = 1.4695$ and $n_{\text{Nb2O5}} = 2.2844$, respectively [77], which is similar in terms of the constituent materials to the coating applied to the silicon nitride membrane to form the membrane mirror (see Chapter 3). In this case, the reflectance already exceeds $\mathcal{R}$QWOT $\gtrsim 0.88$ for $N = 8$ layers, and increases to $\mathcal{R}_{\text{QWOT}} \gtrsim 0.997$ for $N = 16$ layers.

The corresponding transmission spectra for stacks with $N = \{8, 16, 24\}$ are shown

in Figure A.1. They illustrate the formation of a high-reflectance (stop) band centred around the design wavelength (here, $\lambda_0 = 780$ nm). The strong spectral selectivity of the structure is evident from the rapid increase in transmittance as the wavelength deviates from λ_0, highlighting the bandwidth-limited nature of the Bragg mirror's reflectivity.

A.2. Intensity distribution within the mirror coating

As discussed in Section 3.3, the silicon nitride layer in the membrane mirror exhibits non-negligible absorption of the incident optical power. The resulting absorptive loss depends sensitively on the local optical intensity inside the multilayer stack. The transfer-matrix formalism provides a convenient way to compute this spatial intensity distribution and thereby quantify absorption in specific layers.

For illumination from the left, no incident field is present on the right-hand side of the stack. Imposing this boundary condition allows the complex reflection amplitude to be obtained directly from the system matrix:

$$\begin{pmatrix} \tau_{\mathcal{S}} \\ 0 \end{pmatrix} = \mathcal{S} \cdot \begin{pmatrix} 1 \\ \rho_{\mathcal{S}} \end{pmatrix} \Rightarrow \rho_{\mathcal{S}} = -\frac{\mathcal{S}_{21}}{\mathcal{S}_{22}}. \tag{A.7}$$

Here, $\mathcal{S}$ denotes the transfer matrix of the full membrane-mirror stack, including the silicon nitride layer (see Section 3.2 for a detailed discussion of the coating composition). This defines the input field vector as $\boldsymbol{a}\mathrm{in} = (1, \rho_{\mathcal{S}})^{\mathrm{T}}$. The field at an internal position d is then obtained by propagating $\boldsymbol{a}_{\mathrm{in}}$ with the product of transfer matrices up to depth d. The corresponding intensity profile follows from the squared magnitude of the resulting local field amplitude.

Figure A.2 shows the computed intensity distribution in the membrane mirror at $\lambda_0 = 780$ nm for illumination from the left (green) and from the right (blue), corresponding to reflection from the coated side and from the silicon nitride side, respectively. The approximately constant standing-wave amplitude within the silicon nitride layer for right-side illumination indicates that the SiN layer participates as a (lossy) intracavity medium, whose local intensity directly follows the intracavity field amplitude. Its influence on the optical response is discussed in Chapter 6.

Parity of the hybridised $|\pm\rangle$ states of the coupled system

In a coupled-cavity configuration, the simplifying boundary condition used above, namely that the optical field vanishes on one side of the stack, is no longer applicable, since the membrane mirror is embedded between two cavities and is driven from both sides by the intracavity fields. Nevertheless, the multilayer stack can still be treated as a single optical component with system matrix $\mathcal{S}$ by embedding it into the full transfer-matrix model of the coupled cavities (i.e. by adding the remaining cavity propagation and mirror matrices on either side). The resulting compound system (say, described by a composite matrix $\mathcal{C}$) has a well-defined complex reflection amplitude $\rho_{\mathcal{C}}$, such that the input field vector can again be written as $\boldsymbol{a}_{\mathrm{in}} = (1, \rho_{\mathcal{C}})^{\mathrm{T}}$.

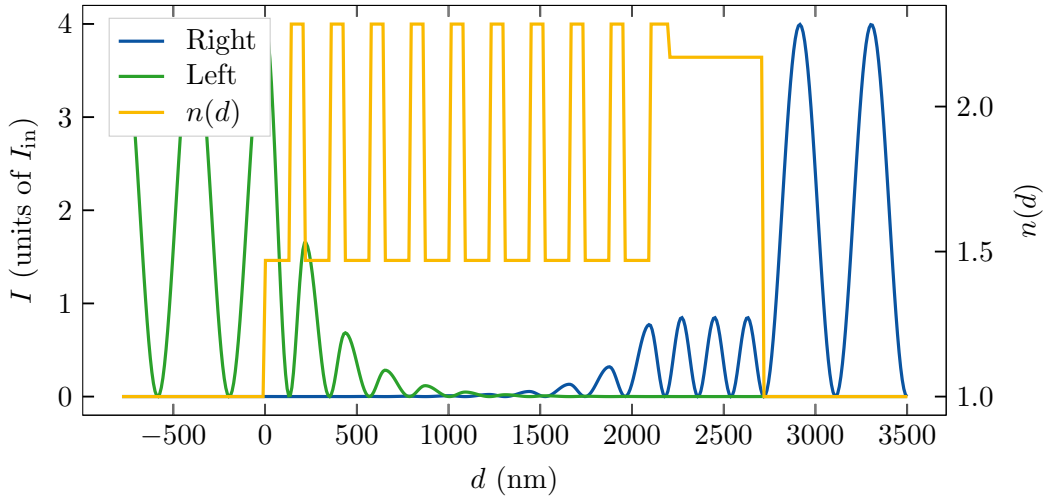


Figure A.2.: Spatial intensity distribution for a field at $\lambda_0 = 780$ nm incident on the membrane mirror from the coated side (left) and from the silicon nitride side (right). The refractive-index profile $n(d)$ is shown on the right axis; the silicon nitride layer has $n_{\mathrm{SiN}} = 2.17$. For illumination from the silicon nitride side, the intensity within the SiN layer is proportional to I_{in}.

To obtain the field inside the structure, $\boldsymbol{a}_{\mathrm{in}}$ is first propagated through the input mirror and the corresponding cavity propagation section. The resulting field vector then serves as the incident field on $\mathcal{S}$, allowing the spatial field (and intensity) distribution inside the membrane mirror to be evaluated in direct analogy to the single-sided case.

Figure A.3 shows the intensity distributions within the membrane mirror for the $|\pm\rangle$ eigenmodes of the coupled-cavity system at the operating point of minimal mode splitting, i.e. at the degeneracy point, see Section 4.1. The intensity profile inside the silicon nitride layer is essentially identical for both states. Only near the centre of the mirror a small difference becomes visible, which is consistent with the opposite parities of the two coupled modes.

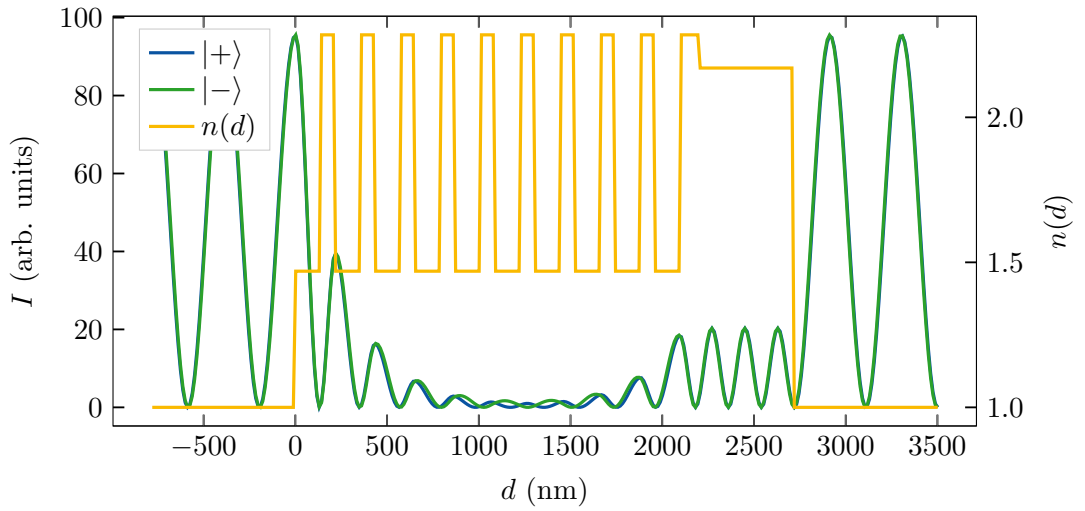


Figure A.3.: Spatial intensity distribution of the coupled-cavity eigenstates $|\pm\rangle$ at $\Delta_{\mathrm{L}} = 0$ and $\lambda \approx 780$ nm, highlighting their opposite parity. Away from the mirror centre, the two intensity profiles become indistinguishable. The right axis indicates the refractive index n as a function of depth into the stack, d. The silicon nitride layer is located on the right-hand side of the structure and has refractive index $n_{\mathrm{SiN}} = 2.17$.

Bibliography

[1] H.-P. Breuer and F. Petruccione. “The Theory of Open Quantum Systems” (Oxford Univ. Press, 2007).

[2] D. A. Lidar. “Lecture Notes on the Theory of Open Quantum Systems”. arXiv:1902.00967 (2020).

[3] F. Marquardt and A. Püttmann. “Introduction to dissipation and decoherence in quantum systems” (2008).

[4] F. Verstraete, M. M. Wolf, and J. I. Cirac. “Quantum computation and quantum-state engineering driven by dissipation”. *Nat. Phys.* **5**, 633–636 (2009).

[5] M. Müller et al. “Engineered Open Systems and Quantum Simulations with Atoms and Ions”. *Advances in Atomic, Molecular, and Optical Physics* **61**, 1–80 (2012).

[6] K. J. Vahala. “Optical microcavities”. *Nature* **424**, 839–846 (2003).

[7] C. W. Gardiner and P. Zoller. “Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics” (Springer, 2004).

[8] W. Bogaerts et al. “Silicon microring resonators”. *Laser & Photon. Rev.* **6**, 47–73 (2012).

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

[10] H. Pfeifer et al. “Achievements and perspectives of optical fiber Fabry–Perot cavities”. *Appl. Phys. B* **128**, 29 (2022).

[11] D. Marcuse. “Coupled mode theory of optical resonant cavities”. *IEEE Journal of Quantum Electronics* **21**, 1819–1826 (1985).

[12] S. Fan, W. Suh, and J. D. Joannopoulos. “Temporal coupled-mode theory for the Fano resonance in optical resonators”. *J. Opt. Soc. Am. A* **20**, 569–572 (2003).

[13] P. Chak, J. K. S. Poon, and A. Yariv. “Optical bright and dark states in side-coupled resonator structures”. *Opt. Lett.* **32**, 1785–1787 (2007).

[14] D. D. Smith et al. “Coupled-resonator-induced transparency”. *Phys. Rev. A* **69**, 063804 (2004).

[15] S. V. Boriskina. “Photonic Molecules and Spectral Engineering”. *Springer Series in Optical Sciences* **156**, 393–421 (2010).

[16] K. Liao et al. "Photonic molecule quantum optics". *Adv. Opt. Photon.* **12**, 60–134 (2020).

[17] R. El-Ganainy et al. "Non-Hermitian physics and PT symmetry". *Nat. Phys.* **14**, 11–19 (2018).

[18] B. E. Little et al. "Microring resonator channel dropping filters". *Journal of Lightwave Technology* **15**, 998–1005 (1997).

[19] Q. Xu et al. "Experimental Realization of an On-Chip All-Optical Analogue to Electromagnetically Induced Transparency". *Phys. Rev. Lett.* **96**, 123901 (2006).

[20] C. Stambaugh et al. "From membrane-in-the-middle to mirror-in-the-middle with a high-reflectivity sub-wavelength grating". *Annalen der Physik* **527**, 81–88 (2015).

[21] N. E. Flowers-Jacobs et al. "Fiber-cavity-based optomechanical device". *Appl. Phys. Lett.* **101**, 221109 (2012).

[22] C. Gärtner et al. "Integrated Optomechanical Arrays of Two High Reflectivity SiN Membranes". *Nano Letters* **18**, 7171–7175 (2018).

[23] J. Xu et al. "Millimeter-scale ultrathin suspended metasurface integrated high-finesse optomechanical cavity". *Opt. Lett.* **47**, 5481–5484 (2022).

[24] H. R. Böhm et al. "High reflectivity high-Q micromechanical Bragg mirror". *Appl. Phys. Lett.* **89**, 223101 (2006).

[25] S. Gröblacher et al. "Demonstration of an ultracold micro-optomechanical oscillator in a cryogenic cavity". *Nat. Phys.* **5**, 485–488 (2009).

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

[27] M. Steiner et al. "Single Ion Coupled to an Optical Fiber Cavity". *Phys. Rev. Lett.* **110**, 043003 (2013).

[28] Y. Colombe et al. "Strong atom–field coupling for Bose–Einstein condensates in an optical cavity on a chip". *Nature* **450**, 272–276 (2007).

[29] F. Rönchen, T. F. Langerfeld, and M. Köhl. "Correlated photon-pair generation in a liquid-filled microcavity". *New J. Phys.* **21**, 123037 (2019).

[30] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt. "Cavity optomechanics". *Rev. Mod. Phys.* **86**, 1391–1452 (2014).

[31] A. İmamoğlu and H. Schmidt and G. Woods and M. Deutsch. "Strongly Interacting Photons in a Nonlinear Cavity". *Phys. Rev. Lett.* **79**, 1467–1470 (1997).

[32] M. J. Hartmann, F. G. S. L. Brandão, and M. B. Plenio. "Strongly interacting polaritons in coupled arrays of cavities". *Nat. Phys.* **2**, 849–855 (2006).

[33] M. J. Hartmann. "Quantum simulation with interacting photons". *J. Opt.* **18**, 104005 (2016).

[34] T. Liew and V. Savona. "Single Photons from Coupled Quantum Modes". *Phys. Rev. Lett.* **104**, 183601 (2010).

[35] H. Flayac and V. Savona. "Single photons from dissipation in coupled cavities". *Phys. Rev. A* **94**, 013815 (2016).

[36] H. Flayac and V. Savona. "Unconventional photon blockade". *Phys. Rev. A* **96**, 053810 (2017).

[37] M. Bamba et al. "Origin of strong photon antibunching in weakly nonlinear photonic molecules". *Phys. Rev. A* **83**, 021802 (2011).

[38] S. Wehner, D. Elkouss, and R. Hanson. "Quantum internet: A vision for the road ahead". *Science* **362**, eaam9288 (2018).

[39] C. Noh et al. "Out-of-Equilibrium Physics in Driven Dissipative Photonic Resonator Arrays". *Quantum Simulations with Photons and Polaritons*, 43–70 (Springer, 2017).

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

[41] S. J. Byrnes. "Multilayer optical calculations". arXiv:1603.02720v5 (2020).

[42] G. Brooker. "Modern Classical Optics" (Oxford Univ. Press, 2006).

[43] B. E. A. Saleh and M. C. Teich. "Fundamentals of photonics: Second Edition" (Wiley, 2007).

[44] S. Hooker and C. Webb. "Laser Physics" (Oxford Univ. Press, 2010).

[45] H. Kogelnik and T. Li. "Laser beams and resonators". *Appl. Opt.* **5**, 1550–1567 (1966).

[46] C. J. Foot. "Atomic physics" (Oxford Univ. Press, 2005).

[47] P. Kobel, R. A. Berner, and M. Köhl. "Secure distribution of a certified random quantum key using an entangled memory qubit". arXiv:2111.14523v1 (2021).

[48] P. Kobel, M. Breyer, and M. Köhl. "Deterministic spin-photon entanglement from a trapped ion in a fiber Fabry–Perot cavity". *npj Quantum Inf.* **7**, 6 (2021).

[49] J. Gallego et al. "High-finesse fiber Fabry–Perot cavities: stabilization and mode matching analysis". *Appl. Phys. B* **122**, 47 (2016).

[50] C. Saavedra et al. "Tunable fiber Fabry-Perot cavities with high passive stability". *Opt. Express* **29**, 974–982 (2021).

[51] P. Fürtjes. "Integrated Fibre Cavity for Quantum Networks" (Master's thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, 2018).

[52] S. Surendra. "Towards quantum non-linear optics with an atomically thin material and optical fibres" (Master's thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, 2020).

[53] M. Scharfstädt. "Implementation of a mircro-mirror production setup for quantum information applications" (Master's thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, 2022).

[54] A. Franzen. "ComponentLibrary". Licensed under Creative Commons Attribution-NonCommercial 3.0 Unported License (CC BY-NC 3.0) (2006).

[55] Metal-coated single-mode fibres (Cu800) data sheet (IVG Fiber, Toronto).

[56] C. Qin et al. "Optical characterization of a fiber Fabry-Perot cavity: precision measurement of intra-cavity loss, transmittance, and reflectance". *Opt. Express* **32**, 14780–14788 (2024).

[57] R. Hauck, H. P. Kortz, and H. Weber. "Misalignment sensitivity of optical resonators". *Appl. Opt.* **19**, 598–601 (1980).

[58] G. Abadias et al. "Review Article: Stress in thin films and coatings: Current status, challenges, and prospects". *J. Vac. Sci. Technol. A* **36**, 020801 (2018).

[59] H. Gibbs. "Optical bistability: Controlling light with light" (Elsevier, 2012).

[60] T. Carmon, L. Yang, and K. J. Vahala. "Dynamical thermal behavior and thermal self-stability of microcavities". *Opt. Express* **12**, 4742–4750 (2004).

[61] L. F. Buchmann et al. "Macroscopic Tunneling of a Membrane in an Optomechanical Double-Well Potential". *Phys. Rev. Lett.* **108**, 210403 (2012).

[62] Y. Li, L.-A. Wu, and Z. D. Wang. "Fast ground-state cooling of mechanical resonators with time-dependent optical cavities". *Phys. Rev. A* **83**, 043804 (2011).

[63] S. Gohlke. "Spectroscopy of coupled cavities" (Master's thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, 2019).

[64] A. M. Jayich et al. "Dispersive optomechanics: a membrane inside a cavity". *New J. Phys.* **10**, 095008 (2008).

[65] B. M. Zwickl et al. "High quality mechanical and optical properties of commercial silicon nitride membranes". *Appl. Phys. Lett.* **92**, 103125 (2008).

[66] X-Ray microscopy windows specification sheet (Norcada Inc., Edmonton).

[67] F. Guo et al. "Effect of thermal annealing on stress relaxation and crystallisation of ion beam sputtered amorphous Si1-xGex thin films". *Mater. Des.* **156**, 389–397 (2018).

[68] W. Ren et al. "Annealing effects on the optical and electrochemical properties of tantalum pentoxide films". *J. Adv. Ceram.* **10**, 704–713 (2021).

[69] D. Wu et al. "Performance Analysis of Surface Reconstruction Algorithms in Vertical Scanning Interferometry Based on Coherence Envelope Detection". *Micromachines* **12**, 164 (2021).

[70] D. Wu and F. Fang. "Development of surface reconstruction algorithms for optical interferometric measurement". *Front. Mech. Eng.* **16**, 1–31 (2021).

[71] C. Yang and J. Pham. "Characteristic Study of Silicon Nitride Films Deposited by LPCVD and PECVD". *Silicon* **10**, 2561–2567 (2018).

[72] M. Ardigo, M. Ahmed, and A. Besnard. "Stoney Formula: Investigation of Curvature Measurements by Optical Profilometer". *Advanced Materials Research* **996**, 361–366 (2014).

[73] A. W. Leissa. "Vibration of Plates" (Government Printing Office, Washington D.C., 1969).

[74] H. L. Offereins, K. Kühl, and H. Sandmaier. "Methods for the fabrication of convex corners in anisotropic etching of (100) silicon in aqueous KOH". *Sensors and Actuators A: Physical* **25**, 9–13 (1990).

[75] D. de Lima Monteiro et al. "Single-mask microfabrication of aspherical optics using KOH anisotropic etching of Si". *Opt. Express* **11**, 2244–2252 (2003).

[76] M. Ghaderi and R. F. Wolffenbuttel. "Design and fabrication of ultrathin silicon-nitride membranes for use in UV-visible airgap-based MEMS optical filters". *J. Phys.: Conf. Ser.* **757**, 012032 (2016).

[77] M. Polyanskiy. "Refractive index data base". URL: https://refractiveindex.info/ (visited on 09/14/2023).

[78] T. F. Langerfeld, H. M. Meyer, and M. Köhl. "Correlated-photon-pair emission from a cw-pumped Fabry-Perot microcavity". *Phys. Rev. A* **97**, 023822 (2018).

[79] T. P. Purdy, R. W. Peterson, and C. A. Regal. "Observation of Radiation Pressure Shot Noise on a Macroscopic Object". *Science* **339**, 801–804 (2013).

[80] J. D. Thompson et al. "Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane". *Nature* **452**, 72–75 (2008).

[81] S. Weis et al. "Optomechanically Induced Transparency". *Science* **330**, 1520–1523 (2010).

[82] T. Wah. "Vibration of Circular Plates". *J. Acoust. Soc. Am.* **34**, 275–281 (1962).

[83] R. W. P. Drever et al. "Laser phase and frequency stabilization using optical resonator". *Appl. Phys. B* **31**, 97–105 (1983).

[84] E. D. Black. "An introduction to Pound–Drever–Hall laser frequency stabilization". *Am. J. Phys.* **69**, 79–87 (2001).

[85] R. F. Voss and J. Clarke. "Flicker ($\frac{1}{f}$) noise: Equilibrium temperature and resistance fluctuations". *Phys. Rev. B* **13**, 556–573 (1976).

[86] S. S. Verbridge, H. G. Craighead, and J. M. Parpia. "A megahertz nanomechanical resonator with room temperature quality factor over a million". *Appl. Phys. Lett.* **92**, 013112 (2008).

[87] V. Jagota, A. Sethi, and D.-K. Kumar. "Finite Element Method: An Overview". *Walailak Journal of Science & Technology (WJST)* **10**, 1–8 (2013).

[88] G. Kirchhoff. "Über das Gleichgewicht und die Bewegung einer elastischen Scheibe." *Journal für die reine und angewandte Mathematik* **40**, 51–88 (1850).

[89] A. E. H. Love. "The Small Free Vibrations and Deformation of a Thin Elastic Shell". *Philosophical Transactions of the Royal Society of London. A* **179**, 491–546 (1888).

[90] A. Khan, J. Philip, and P. Hess. "Young's modulus of silicon nitride used in scanning force microscope cantilevers". *J. Appl. Phys.* **95**, 1667–1672 (2004).

[91] E. Çetinörgü et al. "Mechanical and thermoelastic characteristics of optical thin films deposited by dual ion beam sputtering". *Appl. Opt.* **48**, 4536–4544 (2009).

[92] R. J. Thompson, G. Rempe, and H. J. Kimble. "Observation of normal-mode splitting for an atom in an optical cavity". *Phys. Rev. Lett.* **68**, 1132–1135 (1992).

[93] M. O. Scully and M. S. Zubairy. "Quantum Optics" (Cambridge Univ. Press, 1997).

[94] J. Basdevant and J. Dalibard. "Quantum Mechanics" (Springer, 2006).

[95] Z. H. Wang, X.-W. Xu, and Y. Li. "Partially dark optical molecule via phase control". *Phys. Rev. A* **95**, 013815 (2017).

[96] R. Shankar. "Principles of quantum mechanics" (Springer, 1980).

[97] L. Mandel and E. Wolf. "Optical Coherence and Quantum Optics" (Cambridge Univ. Press, 1995).

[98] S. C. Doret. "Simple, low-noise piezo driver with feed-forward for broad tuning of external cavity diode lasers". *Rev. Sci. Instrum.* **89**, 023102 (2018).

[99] A. Sannia et al. "Dissipation as a resource for Quantum Reservoir Computing". *Quantum* **8**, 1291 (2024).

[100] E. M. Purcell. "Spontaneous Emission Probabilities at Radio Frequencies". *Phys. Rev.* **69**, 681 (1946).

[101] K.-J. Boller, A. Imamoğlu, and S. E. Harris. "Observation of electromagnetically induced transparency". *Phys. Rev. Lett.* **66**, 2593–2596 (1991).

[102] M. Fox. "Quantum optics: an introduction" (Oxford Univ. Press, 2006).

[103] A. E. Krasnok et al. "An antenna model for the Purcell effect". *Sci. Rep.* **5**, 12956 (2015).

[104] S.-T. Guo et al. "Transition between coupled-resonator-induced transparency and absorption". *Phys. Rev. A* **103**, 033510 (2021).

[105] B. Peng et al. "What is and what is not electromagnetically induced transparency in whispering-gallery microcavities". *Nat. Comm.* **5**, 5082 (2014).

[106] R. Finkelstein et al. "A practical guide to electromagnetically induced transparency in atomic vapor". *New J. Phys.* **25**, 035001 (2023).

[107] C. Liu et al. "Observation of coherent optical information storage in an atomic medium using halted light pulses". *Nature* **409**, 490–493 (2001).

[108] K. Totsuka, N. Kobayashi, and M. Tomita. "Slow Light in Coupled-Resonator-Induced Transparency". *Phys. Rev. Lett.* **98**, 213904 (2007).

[109] C.-L. Tien and T.-W. Lin. "Thermal expansion coefficient and thermomechanical properties of SiNx thin films prepared by plasma-enhanced chemical vapor deposition". *Appl. Opt.* **51**, 7229–7235 (2012).

[110] H. Nejadriahi et al. "Thermo-optic properties of silicon-rich silicon nitride for on-chip applications". *Opt. Express* **28**, 24951–24960 (2020).

[111] C. Li. "All-Optical Switch Based on Nonlinear Optics". *Nonlinear Optics: Principles and Applications*, 279–386 (Springer, 2017).

[112] A. M. C. Dawes et al. "All-Optical Switching in Rubidium Vapor". *Science* **308**, 672–674 (2005).

[113] M. D. Lukin. "Colloquium: Trapping and manipulating photon states in atomic ensembles". *Rev. Mod. Phys.* **75**, 457–472 (2003).

[114] V. Scarani et al. "The security of practical quantum key distribution". *Rev. Mod. Phys.* **81**, 1301–1350 (2009).

[115] K. M. Birnbaum et al. "Photon blockade in an optical cavity with one trapped atom". *Nature* **436**, 87–90 (2005).

[116] K. Mukherjee and P. C. Jana. "Optical Bistability in Coupled Cavity System". *ICOL-2019 Proceedings*, 247–250 (2021).

[117] F. Zou, D.-G. Lai, and J.-Q. Liao. "Enhancement of photon blockade effect via quantum interference". *Opt. Express* **28**, 16175–16190 (2020).

[118] J. M. P. Nair, D. Mukhopadhyay, and G. S. Agarwal. "Enhanced Sensing of Weak Anharmonicities through Coherences in Dissipatively Coupled Anti-PT Symmetric Systems". *Phys. Rev. Lett.* **126**, 180401 (2021).

[119] H. Pan et al. "Bistability in dissipatively coupled cavity magnonics". *Phys. Rev. B* **106**, 054425 (2022).

[120] D. Milam. "Review and assessment of measured values of the nonlinear refractive-index coefficient of fused silica". *Appl. Opt.* **37**, 546–550 (1998).

[121] M.-C. Tien et al. "Ultra-low loss Si3N4 waveguides with low nonlinearity and high power handling capability". *Opt. Express* **18**, 23562–23568 (2010).

[122] C.-Y. Tai et al. "Determination of nonlinear refractive index in a Ta2O5 rib waveguide using self-phase modulation". *Opt. Express* **12**, 5110–5116 (2004).

[123] R. Conti, A. A. Gallitto, and E. Fiordilino. "Measurement of the Convective Heat-Transfer Coefficient". *Phys. Teach.* **52**, 109–111 (2014).

[124] W. Chen et al. "A simple method for characterizing and engineering thermal relaxation of an optical microcavity". *Appl. Phys. Lett.* **109**, 061103 (2016).

[125] H. Ftouni et al. "Thermal conductivity of silicon nitride membranes is not sensitive to stress". *Phys. Rev. B* **92**, 125439 (2015).

[126] T. Gu et al. "Molecular-absorption-induced thermal bistability in PECVD silicon nitride microring resonators". *Opt. Express* **22**, 18412–18420 (2014).

[127] C. Qiu et al. "All-optical control of light on a graphene-on-silicon nitride chip using thermo-optic effect". *Sci. Rep.* **7**, 17046 (2017).

[128] Y. Wang et al. "Quality factor measurement for MEMS resonator using time-domain amplitude decaying method". *Microsyst. Technol.* **21** (2014).

[129] F. Zhou et al. "Cavity optomechanical bistability with an ultrahigh reflectivity photonic crystal membrane". *Laser & Photon. Rev.* **17**, 2300008 (2022).

[130] N. Vukovic et al. "Ultrafast optical control using the Kerr nonlinearity in hydrogenated amorphous silicon microcylindrical resonators". *Sci. Rep.* **3**, 2885 (2013).

[131] P. Zhao et al. "Temporal and polarization dependence of the nonlinear optical response of solvents". *Optica* **5**, 583–594 (2018).

[132] K. Kuruma et al. "Time-resolved vacuum Rabi oscillations in a quantum-dot–nanocavity system". *Phys. Rev. B* **97**, 235448 (2018).

[133] C. Vaneph et al. "Observation of the Unconventional Photon Blockade in the Microwave Domain". *Phys. Rev. Lett.* **121**, 043602 (2018).

[134] S. R. K. Rodriguez et al. "Interaction-induced hopping phase in driven-dissipative coupled photonic microcavities". *Nat. Comm.* **7**, 11887 (2016).

Acknowledgements

First and foremost, I would like to express my sincere gratitude to Prof. Dr. Michael Köhl, not only for giving me the opportunity to work on this project, but also for his guidance, support, and trust throughout the entire journey.

I am also grateful to Eric Daigle (Norcada Inc.) and Tobias Groß (Laseroptik GmbH) for their essential contributions to the membrane mirror prototype that made this thesis possible in the first place. Their work on the fabrication of tailored silicon nitride membranes and the optimisation of the applied dielectric coating, respectively, together with many lengthy discussions, were crucial for bringing the prototype to life. Furthermore, I would like to thank the colleagues from the institute's workshop, in particular Alexander Christen, for their great input and the fast realisation of custom assemblies that formed the backbone of this experiment.

Of course, I also want to thank all members of the "Quantum group", past and present, who crossed my path along the way. In particular, I would like to emphasise Thorsten, Felix, and Andrea, not only for proofreading parts of this thesis, but even more for their support beyond that:

Thorsten, thank you for being there throughout my journey. You believed in me even when I did not, and you always seemed to know exactly what to say to give me confidence. I always looked up to you, and I greatly appreciated having you as a mentor. I am especially grateful for our discussions leading up to the membrane mirror prototype; they had a significant impact on the final design and ultimately on the success of this thesis.

Felix, I am grateful that we truly got to know each other during our master's theses, even though we started studying in the same year. Sometimes I wish we had become friends earlier, but I am even more glad that we did when we did. I look back fondly on the summer school in Innsbruck and the time we spent there. More than that, you were one of the main drivers pushing me to finish this thesis, especially when my energy and motivation were running low. Thank you for your persistence, your honesty, and for always being in my corner.

Thorsten and Felix, I am unendingly grateful for the day you cornered me in the office and confronted me. For me, that moment was a turning point. I know I tested your patience in the months after that, but in the end I made it because of that day, and I will never forget it.

Last, but definitely not least, Andrea, thank you as well. When I started to doubt the experiment and felt my spark fading, you helped me reignite it. Our discussions shaped not only the direction of the work, but also the way I approached the thesis.

Overall, I am immensely grateful that you three were a fundamental part of this thesis, and, in many ways, of the entire experiment. Thank you.

Moritz, what began as being “imprisoned” in the same lab evolved into a friendship, and, by now, into one of the closest friendships I have. I enjoyed every shared tea time and lunch break, and I hope there will be plenty more in the future, wherever life may guide us.

Many thanks also go to Akos Hoffmann, whose wisdom in electronics greatly improved the lab experience. In particular, I would like to thank him for the many hours he spent studying, building, and testing an ultrafast electronic pulse generator for the experiment. I hope to implement the device at some point in the future.

I would also like to thank Andreas K., Andreas v. H., Daniel, Janek, Jonas, Luisa, Marcell, Martin, Moritz B., Nick, Nicola, Pascal, Pia, Santhosh, Tobias, and Valentin for their support, whether through borrowed equipment, helpful advice, an open ear, or simply all the shenanigans beyond work. You made this an unforgettable journey and created an environment that truly felt like home.

Finally, I would like to thank my family and friends for their moral support and for accompanying me through all the ups and downs, even when I was not always easy to deal with. In particular, I thank my sister Hanna, who always had my back and found ways to cheer me up no matter how frustrated I was; and my best friend Niklas, who was always there whenever I needed distraction.

Very last, I want to thank my grandparents, Nana and Opa. Without you, I would not have been in the position to start this journey in the first place. It breaks my heart that you are not here anymore to celebrate the end with me.