Imperial College
London

Integrated optics for coupled-cavity quantum electrodynamics

Guillaume Lepert

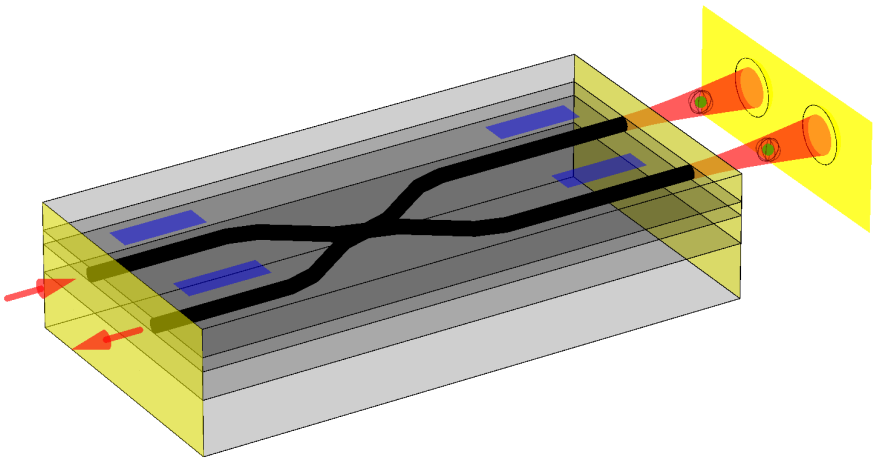

A thesis submitted for the degree of Doctor of Philosophy
of Imperial College London

Physics Department
Quantum Optics and Laser Science/Centre for Cold Matter

January 2013

Abstract

This thesis reports on the development of an array of plane-concave Fabry-Pérot microcavities containing atoms (or other quantum emitters), interconnected by UV-written waveguides on a silica-on-silicon chip. The microcavities are formed by a mirror coated on the end facet of the chip and an array of spherical micromirrors etched on silicon. This is to our knowledge the first attempt at implementing the emerging coupled-cavities QED paradigm. The device we propose possesses a degree of control, flexibility and tuning unmatched in other suggested implementations: The atoms can be manipulated inside the cavity by auxiliary lasers and the cavity-cavity coupling rate as well as the atom-cavity coupling can be tuned. It is highly scalable.

Calculation of the complete (classical) optical spectrum of the device is presented. The quantum dynamics that may eventually be observed has also been studied. Waveguide chips containing couplers and phase shifter have been fabricated. We have successfully demonstrated the operation of the elementary sub-systems: the strong optical coupling between a microcavity and a waveguide resonator, and the tunable strong coupling between two evanescently coupled waveguide resonators.

No experiments with atoms or other quantum emitters were attempted, because the waveguide propagation loss is so large that no quantum physics can be observed. There is hope that this can be overcome in the future by using other waveguide technologies.

Statement Of Originality

I hereby declare that this thesis and the work presented herein are my own, except where otherwise credited.

Referees

Professor Vahid Sandoghdar
Max Planck Institute for the Science of Light
Erlangen

Professor Mark Neil
Physics Department, Photonics Group
Imperial College London

Contents

Acknowledgements

I am forever in debt to Ed, Michael and Jaesuk for the chance of working and learning alongside them. I have had many enlightening discussions with and received invaluable help from the bright people at CCM; Rob, Adrian, Joe and Stefan deserve special thanks. Kudos to Steve, Jon and Val in the workshop: nothing is impossible for them. Special mention goes to Helen and James in Southampton, with whom it has been a real pleasure to collaborate.

It's been (mostly) fun.

Chapter 1

A new dawn for cavity QED

> The interaction of a single dipole with a monochromatic radiation field presents an important theoretical problem in electrodynamics. It is an unrealistic problem in the sense that experiments are not done with single atoms and single-mode fields.
>
> — L Allen and JH Eberly, *Optical resonance and two-level atoms*, 1975 (p. 157)

> The Nobel Prize in Physics 2012 was awarded jointly to Serge Haroche and David J. Wineland "for ground-breaking experimental methods that enable measuring and manipulation of individual quantum systems".
>
> — The Royal Swedish Academy of Sciences, 9 October 2012

1.1 The problem with atomic physics

The two quotes above illustrate in a rather dramatic way the huge progress that has been made in the field of light-matter interactions in the last decades. For Allen and Eberly, and for all atomic physicists at the time, the semiclassical approximation was sufficient for all practical purpose; quantum electrodynamics was merely an intellectual curiosity. Yet shortly before the Nobel prize announcement Haroche demonstrated full quantum feedback control, using single atoms to inject or remove – on demand – single photons from the electromagnetic field inside a cavity (1).

The traditional difficulty with atomic physics is two-fold: the light-matter interaction is weak, and the probability of detecting photons is very low. As a result, experiments demanded that many photons scatter off many atoms, and the realm of interactions between single quanta remained elusive.

It is not hard to see why this had to be. The collection efficiency of a standard optical system, as a function of its numerical aperture NA and assuming isotropic emission, is

$$\eta_{\text{coll}} = \frac{1}{2}\left(1 - \sqrt{1 - \text{NA}^2}\right). \tag{1.1}$$

The NA is defined as NA=$n \sin\theta$, where θ is the optics' collection angle and n the refractive index at the source. Thus a low-NA optics observing an atomic ensemble from outside the vacuum chamber would only collect $\mathrm{NA}^2/4$ of all emitted photons — typically as low as 10^{-4}. Even a large NA system, say an NA=0.5 diffraction-limited asphere lens, placed inside the vacuum, only boosts η_{coll} to a 5%.

On the other hand, the absorption cross-section of a simple two-level atom

$$\sigma_{\mathrm{abs}} = \frac{3\lambda^2}{2\pi} \tag{1.2}$$

is typically an order of magnitude or two smaller than the focussed spot size achievable with such an optical system, and so the interaction probability is very small. For the sake of concreteness, let's consider a single rubidium atom and calculate, in a very crude semiclassical way, the probability that a single 780 nm photon may excite it to the upper state of the D2 transition. We assume that the photon wavepacket has a duration $\tau = 3\,\mathrm{ns}$ (which is much shorther than the D2 spontaneous lifetime $\gamma \sim$MHz) with a top hat profile and a $50\,\mu$m waist. The Rabi frequency is then $\frac{g}{2\pi} \sim 0.2\,\mathrm{MHz}$ and the probability of exciting the atom is only $\sin^2(g\tau/2) \sim 3 \times 10^{-6}$.

The road to seeing single quanta interact is not necessarily closed however, for there are two obvious ways of solving these problems of weak collection and interaction probability. The first one is to match the optical and atomic cross-section, which is in principle possible since the latter is of the same order as the diffraction limited spot size w_{diff}:

$$w_{\mathrm{diff}} \simeq \frac{\lambda}{2\sin\theta} \sim \sqrt{\sigma_{\mathrm{abs}}}, \tag{1.3}$$

λ being the wavelength in the medium. The second possibility is to confine a photon between two mirrors so that it interacts with the same atom repeatedly. The next section will briefly outline recent progress along the first path. We'll then move on to cavity QED, which is the focus of this work.

1.2 Strong light-matter interactions I: Shrinking the photons

1.2.1 Solid immersion lenses

In order to achieve high efficiency interaction between an emitter and light, the spatio-temporal overlap between the photonic mode and the dipole radiation pattern must be as high as possible. It has been shown theoretically that a "dipole wave" could be perfectly reflected by an oscillating dipole (2; 3). It is possible but difficult to produce such a wave (4); however it is reasonably approximated by a strongly focused Gaussian beam,

for which a reflection by a single dipole of up to 85% has been predicted (2). In order to achieve the required focusing, Sandoghdar used "solid immersion lenses", similar to high-NA oil immersion objectives except that the index-matching oil has been replaced by a solid hemispherical lens coated with an organic crystalline matrix. Dye molecules such as DBATT or DBT are embedded in the matrix; at 1.5 K the zero-phonon line (ZPL) of these molecules become lifetime-limited so that it behaves as a nearly perfect two-level-system. A "single-molecule transistor" that could absorb up to 12% of an incoming weak CW probe laser was demonstrated (5); population inversion could be achieved with a second laser and amplification of the probe beam was demonstrated, albeit with a gain of less than one per cent. In a similar experiment the phase shift induced on the probe beam was tuned with a control laser (6).

The collection efficiency of photons emitted by a single molecule can be further boosted by exploiting the modification of the dipole radiation pattern in the thin crystalline film. The proximity of the crystal-air interface strongly redirects the emission away from that interface, and a 96% collection efficiency into an underlying oil immersion objective has been reported (7). The emission frequency of these dye molecules can be tuned by applying an electric field (Stark shift). In this way, two independent molecules could be tuned to the same frequency and a Hong-Ou-Mandel experiment revealed that the emitted photons were indistinguishable (8). Thus these molecules could become a precious resource for quantum information. Sandoghdar also performed a spectroscopy experiment where Stark-shifted single photon emitted by a single molecule were used to probe a second molecule (9).

Solid immersion lens-based systems as well as cryogenic dye molecules are very promising. Cavity enhancement of the molecules' spontaneous emission has also been demonstrated (10), and it may be possible to combine SILs with cavities to increase the nonlinearities as well as the collection efficiencies and, more generally, achieve a finer control of the light-matter interaction. The dye molecules are only two-level systems however, and for many QIP schemes we require three or four level.

1.2.2 High-index nanowaveguides

Nanowires provide an alternative approach to diffraction-limited focusing. These can be thin ($\sim 100\,\mathrm{nm}$) ridge waveguides of a high-index material (for example silicon nitride or tantalum oxide, $n \simeq 2$), or tapered optical fibres with a diameter around half a micron. Either way, the high index contrast between the waveguide and its surroundings results in a high degree of confinement; at the same time the waveguide is so small that the field extends significantly outside, where it can interact with emitters. Experiments with silicon nitride were pioneered by Lipson who used a nanowaveguide chip to probe a hot rubidium vapour (11); work on silicon nitride and dye molecules is ongoing at the Centre

for Cold Matter (CCM) (12).

Most of the work in this area has relied on nanofibres. They can be used to probe a cold atomic ensemble (13), and as traps for neutral atom (14). The photonic mode density is strongly altered in the vicinity of the nanofibre, so that the emission is strongly (up to 25 per cent) redirected into the guided modes (15) (see also the calculation in Appendix B). This can be boosted to 90% by defining a Bragg grating cavity directly on the nanofibre (16) or in the non-tapered region (17).

In addition there are other systems with similar features. For example a diamond sample containing nitrogen vacancy (NV) centres can be etched into a nanowaveguide ring cavity to enhance the ZPL emission (18). Even tighter light confinement can be achieved in plasmonic systems, but at the cost of large propagation losses.

Last but not least, we should mention that instead of tightly focusing the light mode to achieve the highest possible value of the electric field, we can inflate the atomic absorption cross-section and the atom dipolar moment by exciting atoms to a Rydberg state. Blocade radii in excess of $10\,\mu$m are possible, and we can use EIT-like schemes to generate strong non-linearities as in Lukin's recent experiment (19).

1.3 Strong light-matter interactions II: Recycling the photons with cavities

Optical cavities enhance the interaction between atoms and photons through the interplay of several effects. First, photons can be confined to a small volume defined by the mirrors and the optical (Gaussian) modes they support; the mode volume can be several orders of magnitude smaller than in free space (excluding ultra-short femtosecond pulses, which we do not consider here since they usually contain a very large number of photons). The Rabi frequency per photon is thus much larger, and in small cavities a single photon can be sufficient to saturate the atomic transition. Second, photons can remain inside the cavity for very long times (up to $\sim 100\,$ms) and have a long time to interact with the atoms. Third, the cavity singles out a particular mode of the electromagnetic field into which atoms decay preferentially, and this mode can be efficiently channelled to a detector or to another experiment.

On the other hand, cavities are experimentally challenging, both to fabricate and to operate. High-reflectivity mirrors are expensive to fabricate, requiring huge expertise and state-of-the-art deposition facilities. High-finesse cavities suffer from a narrow bandwidth and therefore need to be tunable and extremely stable (although only on relatively short time scales, given that experiments rarely last more than a few milliseconds). This also makes it difficult to integrate more than one type of emitters in the cavity.

In the following section we describe the main features of cavity quantum electrody-

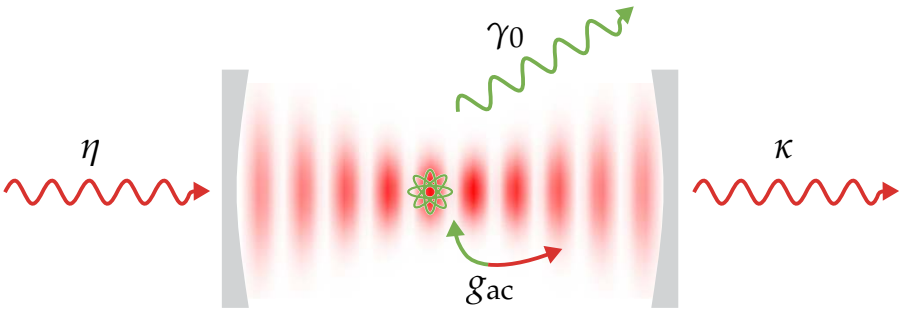

Figure 1.1: An atom in a cavity, showing the atom-cavity Rabi frequency g_{ac}, the atom free space spontaneous emission γ_0, the cavity decay rate (half-linewidth) κ and the cavity pumping rate η.

namics in the Jaynes-Cumming model. We then review its historical development, with a focus on the very early works of Purcell and Drexhage.

1.3.1 Fundamentals of cavity QED

The theoretical framework for cavity QED was laid out by Jaynes and Cummings in 1963 (20). They established the full quantum electrodynamic theory of a single emitter (in this case, an ammonia molecule in a maser) interacting with a single mode of the electromagnetic field. The paradigm has not changed much since, except for the inclusion of damping.

The atom-cavity problem is closely related to cooperative phenomena in atomic ensembles, such as Dicke's superradiance (21). Indeed, one can rewrite the atom-cavity problem as a chain of atoms, each of which is a virtual image of the initial atom, formed by the two cavity mirrors (22).

Hamiltonian

Consider the system depicted in Figure 1.1. A single two-level atom with ground and excited states $|g\rangle$ and $|e\rangle$ and a transition energy $\hbar\omega_a$ sits in a cavity, whose resonant frequency is ω_c. The cavity is taken to be a quantum harmonic oscillator, as depicted in Figure 1.2. The atom emits spontaneously at a rate γ and the field intensity in the cavity decays at a rate κ. η describe cavity pumping by an external laser, and g_{ac} is the coupling between the atom and the cavity field.

The cavity field operator is

$$\hat{\mathbf{E}}(r) = \sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}}\left(\psi(\mathbf{r})\hat{a} + \psi^*(\mathbf{r})\hat{a}^\dagger\right)\boldsymbol{\epsilon} \tag{1.4}$$

where $V = \int |\psi(\mathbf{r})|^2 d^3\mathbf{r}$ is the mode volume and $\hat{a}$, $\hat{a}^\dagger$ are the ladder operators for the quantum cavity oscillator. $\boldsymbol{\epsilon}$ describes the polarisation. This definition ensures that the vacuum field energy $\int \varepsilon_0 \langle 0|E(\mathbf{r})^2|0\rangle d^3\mathbf{r}$ has the required value $\hbar\omega/2$.

For a Gaussian Fabry-Pérot mode (see Appendix A), we take

$$\psi(\mathbf{r}) = e^{-(x^2+y^2)/w_0^2} \cos kz. \tag{1.5}$$

$k = 2\pi\lambda$ is the light k-vector in the cavity. On axis and at an antinode, $\psi = 1$.

The coupling g is defined by

$$\hbar g = \boldsymbol{\mu}.|\langle 0|\hat{\mathbf{E}}|1\rangle| = \boldsymbol{\mu}.\boldsymbol{\epsilon}\sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}}|\psi(\mathbf{r})|, \tag{1.6}$$

where $\boldsymbol{\mu}$ is the dipole matrix element. The maximum coupling is therefore

$$g_{\mathrm{ac}} = \sqrt{\frac{\mu^2\omega}{2\hbar\varepsilon_0 V}} \tag{1.7}$$

The Jaynes-Cummings (JC) Hamiltonian takes the form

$$H = \overbrace{\hbar\omega_a|e\rangle\langle e|}^{\text{atom}} + \overbrace{\hbar\omega_c(\hat{n}+1/2)}^{\text{cavity}} + \overbrace{\hbar g_{\mathrm{ac}}\left(|e\rangle\langle g|\hat{a} + \hat{a}^\dagger|g\rangle\langle a|\right)}^{\text{interaction } H_I} \qquad = H_{\mathrm{JC}} \tag{1.8}$$

$$- \underbrace{i\hbar\kappa\hat{a}^\dagger\hat{a}}_{\text{cavity decay}} - \underbrace{i\hbar\gamma_0|e\rangle\langle e|}_{\text{spont. emiss.}} - \underbrace{i\hbar\eta(\hat{a}-\hat{a}^\dagger)}_{\text{cavity pumping}} \qquad = H_{\mathrm{diss}} \tag{1.9}$$

We have introduced the number operator $\hat{n} = \hat{a}\hat{a}^\dagger$. The first part H_{JC} contains only Hermitian (energy-conserving) terms; we have dropped the fast-oscillating and non-energy-conserving terms of the interaction under the usual rotating wave approximation. This then describes the coherent exchange of excitation between the atom and the cavity, and has a block-diagonal structure

$$H_{\mathrm{JC}} = \begin{pmatrix} \frac{1}{2}\omega_c & 0 & 0 & \cdots \\ 0 & H_1 & 0 & \cdots \\ 0 & 0 & H_2 & \cdots \\ \vdots & \vdots & \vdots & H_n \end{pmatrix} \quad \text{with } H_n = \begin{pmatrix} (n+1/2)\omega_c & g\sqrt{n} \\ g\sqrt{n} & (n+1/2)\omega_c + \Delta_{\mathrm{ac}} \end{pmatrix} \tag{1.10}$$

Each block H_n describes a closed two-level subsystem with n excitations in the basis $\{|g,n\rangle, |e,n-1\rangle\}$. $\Delta_{\mathrm{ac}} = \omega_a - \omega_c$ is the atom-cavity detuning. H_n (and therefore H_{JC}) can be diagonalised, and the eigenmodes and eigenfrequencies are

$$|n,\pm\rangle = \frac{1}{\sqrt{2}}\left(|g,n\rangle \pm |e,n-1\rangle\right) \tag{1.11}$$

$$\Omega_n^{\pm} = (n+\frac{1}{2})\omega_c + \frac{1}{2}\left(\Delta_{\mathrm{ac}} \pm \sqrt{4ng_{\mathrm{ac}}^2 + \Delta_{\mathrm{ac}}^2}\right) \tag{1.12}$$

The Jaynes-Cummings energy spectrum is thus a ladder of doublets with increasing sep-

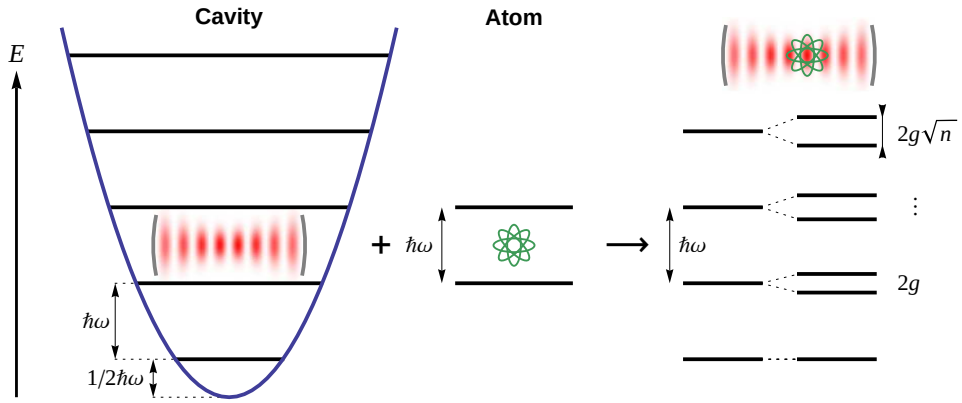

Figure 1.2: Graphical representation of the Jaynes-Cummings Hamiltonian. The cavity field is a harmonic oscillator (left) and the atom a two-level system (centre); the energy spectrum of the coupled system is a ladder of doublets whose separation is the effective Rabi frequency (right).

aration $\Omega_{\text{eff},n} = 2g\sqrt{n}$ (taking $\Delta_{\text{ac}} = 0$), as shown in Figure 1.2. $\Omega_{\text{eff},n}$ is known as the effective Rabi frequency. The system oscillates between the states $|g, n\rangle$ and $|e, n-1\rangle$ at a frequency $\Omega_{\text{eff},n}$; that is, the atom periodically emits and reabsorbs a cavity photon. This is Rabi flopping. Note that the quantum Rabi frequency is no different from its semiclassical counterpart.

The second part H_{diss} of the JC Hamiltonian consists of non-Hermitian terms describing respectively the loss of photons at a rate κ (including, but not limited too, leakage through the cavity mirrors); the loss of the excitation via spontaneous emission outside of the cavity mode, at a rate γ_0; and the injection of coherent states at a rate η, corresponding to laser pumping of the cavity through one of the mirrors[1]. A rigorous justification of these terms would requires us to delve into the intricacies of the master equation formalism; for our purpose it is enough to empirically check that they produce the desired behaviour when injected into the Schrödinger equation.

Consider the damped, non-driven Hamiltonian. When the cavity is resonant with the atom, $\Delta_{\text{ac}} = 0$, the eigenfrequencies are (23)

$$\omega_n^{\pm} = (n + \frac{1}{2})\omega_c \pm \frac{1}{2}\sqrt{4ng_{\text{ac}}^2 - (\gamma_0 - \kappa)^2} - \frac{i}{2}[\gamma_0 + \kappa(2n-1)]. \tag{1.13}$$

Thus the damping not only gives a finite width to the JC doublet resonances, it also changes the Rabi splitting $\omega_n^+ - \omega_n^-$.

We can also calculate the full atom-cavity spectrum that would be observed if one were to probe the cavity with a laser while monitoring the transmission. If the pumping

[1] η is related to the optical power in the input mode P_{in} by $P_{\text{in}} = 2\eta^2\hbar\omega/\kappa$. This is straightforward to prove, starting from the intracavity photon number $n_{\text{ph}} = \eta^2/\kappa^2$. The *power* leaking out from the cavity is then $n_{\text{ph}}\hbar\omega \cdot 2\kappa$, and in the steady-state regime this must be equal to P_{in}.

is weak enough that $n \leq 1$, we can restrict the state vector to

$$|\Psi(t)\rangle = c_e(t)|e,0\rangle + c_g(t)|g,1\rangle + c_0(t)|g,0\rangle \tag{1.14}$$

with $c_0(t) \simeq 1$. The steady-state solution to the Schrödinger equation $i\hbar\dot{\Psi}(t) = H\Psi(t)$ then gives us the population of the cavity field as a function of the laser-cavity detuning Δ_{CL}, assuming that the cavity is tuned to the atomic resonance ($\Delta_{\rm ac} = 0$)

$$\langle n \rangle = |c_g|^2 = \frac{\eta^2}{\kappa^2} \left| \frac{\kappa(\gamma - i\Delta_{CL})}{(\Delta_{CL} - \omega_1^+)(\Delta_{CL} - \omega_1^-)} \right|^2 . \tag{1.15}$$

The cavity transmission is proportional to Eq. (1.15), and is plotted in Figure 1.3(a).

Strong coupling regime

From Eq. (1.13) it is clear that the "vacuum" Rabi splitting can only be resolved if 1) the square root term is real (otherwise the system would be overdamped) and 2) it is larger than the imaginary part, i.e.

$$\begin{cases} 4g_{\rm ac}^2 & > & (\gamma_0 - \kappa)^2 \\ 4g_{\rm ac}^2 & \gtrsim & 2(\gamma_0^2 + \kappa^2) \end{cases} . \tag{1.16}$$

In general we strive for the *strong coupling regime* $g_{\rm ac} \gg \kappa, \gamma_0$. The time evolution of the atomic and photonic populations are then damped sinusoidal oscillations.

Collapse and revival of the Rabi oscillations So far we have assumed that cavity was populated with a number state $|n\rangle$. In practice however, the cavity is pumped with a laser and excitation are fed into the cavity not as Fock states, but as coherent states $|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle$ (a Poisson distribution of Fock states with mean photon number and variance $|\alpha|^2$). Each of these coherent states will undergo Rabi oscillations with the atom, with each Fock component having a different frequency $\Omega_{{\rm eff},n}$. The coherent superposition of these oscillations will eventually kill the Rabi flopping as they dephase, but as long as the coherence is not lost they will eventually rephase and the Rabi oscillations will revive, as shown in Figure 1.3(b).

This collapse and revival effect is qualitatively similar to Newton fringes in white light interferometry and other phenomena involving coherent dephasing, but it is nonetheless a uniquely quantum feature since it is associated with the quantisation of the electromagnetic field. In fact Jaynes and Cummings showed that a semiclassical treatment would reproduce most of the features of cavity QED with reasonable accuracy, with the notable exception of the collapse and revival.

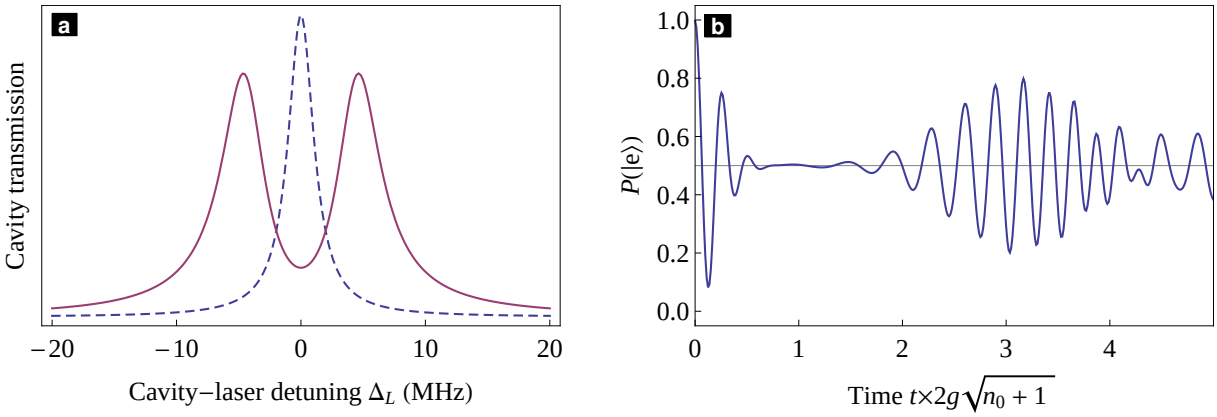

Figure 1.3: (a) Transmission spectrum of a weakly driven cavity containing a single atom (solid line) as a function of the cavity-laser detuning $\Delta_L = \omega_c - \omega_L$, showing the normal-mode splitting. The atom is assumed to resonant with the cavity. The average photon number is $\langle n\rangle = 1$ and the parameters are $g, \gamma, \kappa = 2\pi \times (4.5, 3, 1.25)$ MHz. The dashed line is the empty cavity spectrum. **(b)** Collapse and revival of the Rabi oscillation with a mean photon number $\langle n\rangle = |\alpha|^2 = 5$.

Photon blockade The collapse and revival of the Rabi oscillation, while being a rather spectacular effect, has a rather limited potential in terms of useful applications. A more practical feature of the model is the anharmonicity of the Jaynes-Cummings ladder, since $\Omega_{\text{eff},n} \propto \sqrt{n}$. In the strong coupling regime, this gives rise to the so-called photon blockade effect: if the cavity contains a resonant photon at $\omega_c \pm g_{\text{ac}}$, a second photon cannot enter the cavity since it is not resonant with either of the two-photon lines $\omega_c \pm \sqrt{2} g_{\text{ac}}$. Obviously we require $g_{\text{ac}}(\sqrt{2}-1) \gg \kappa, \gamma$: the one- and two-photon transitions must be clearly resolved. Thus, an incident laser beam at $\omega_c \pm g_{\text{ac}}$ will be converted from a flux of coherent states into a flux of single photons. Indeed, Kimble observed antibunching in the light transmitted through a suitably prepared cavity (24). Conversely, a laser beam at $\omega_c - \sqrt{2} g_{\text{ac}}$ would exhibit bunching. The photon blockade effect is effectively a photon-photon interaction mediated by the cavity-enhanced atom.

Fast cavity regime

Let us now turn to the *fast cavity regime*, $\kappa \gg g$. The photon leaks out of the cavity immediately after being emitted so that the population in $|g, 1\rangle$ can be neglected. Starting with an atom in the excited state $|e, 0\rangle$, we can evolve the damped Hamiltonian (1.9) in time under Schrödinger's equation with the state vector (1.14). The time dependence of the excited state is

$$c_e(t) = e^{-\left(\gamma + \frac{g^2}{\kappa}\right)t} = \exp\left[-\gamma(1+2C)t\right] \quad \text{with} \quad C = \frac{g^2}{2\kappa\gamma}. \tag{1.17}$$

C is known as the cooperativity. This shows how the resonant cavity can enhance the spontaneous emission rate. Interestingly, the extra factor $2C$ is emitted *into the cavity*

mode, and therefore can be efficiently collected. This highlights the potential of cavity QED as a source of single photon. Nonetheless we should keep in mind that photons produced in this way are typically very long ($\sim$ ns), since they cannot be shorter than the cavity lifetime: we are far from producing "light bullets".

Photons of controlled duration and temporal profile (although still limited by the cavity lifetime) can be created if instead of a two level atom we use a three-level Λ-system with two metastable ground states, $|a\rangle$ and $|b\rangle$ coupled to an excited state $|e\rangle$. One of the transitions is coupled to the cavity field, with Rabi frequency $\Omega_c = 2g\sqrt{n}$: $|a, n\rangle \leftrightarrow |e, n-1\rangle$ and the second one to a control laser field Ω_p, with no effect on the cavity photons: $|b, n-1\rangle \leftrightarrow |e, n-1\rangle$. This is analogous to the familiar electromagnetically induced transparency configuration (EIT) (25); in particular one of the three eigenstates is a "dark" state

$$|u\rangle = \cos\Theta|a, n\rangle - \sin\Theta|b, n-1\rangle, \text{ with the mixing angle } \tan\Theta = \frac{\Omega_p}{\Omega_c} \tag{1.18}$$

Thus, by adiabatically sweeping the mixing angle Θ by mean of the control laser Ω_p, it is possible to deterministically create (or destroy) a cavity photon, whose temporal profile can be tailored via the control pulse (26).

1.3.2 Historical overview

NMR and Purcell's seminal "paper"

What is now known as the Purcell effect initially arose in connection with early research in nuclear magnetic resonance, discovered independently by Purcell (27) and Bloch (28) in the winter 1945-46. The principle of NMR was quite well established at the time: protons (hydrogen nuclei) have a spin, whose degeneracy can be lifted in a strong magnetic field, leading to a radiofrequency transition around 30 MHz. The spins align themselves with the magnetic field, initially with equal populations in the "up" and "down" states. Interactions with the environment then bring the spins into thermal equilibrium where the relative populations follow the Boltzmann distribution, which results in a small excess of spin down, of the order of 10^{-5}. Detecting the absorption (of a radiation resonant with the spin transition) induced by this excess requires a large number of spins, which earlier experiments with molecular beams could not provide. Instead Purcell placed a kilogram of solid paraffin inside a metal cavity driven by an RF generator, and placed the apparatus between the poles of a large magnet (the very same that had discovered the muon in 1937) (29). By increasing the current into the magnet, Purcell was able to bring the nuclear spin in resonance with the fixed-frequency RF field, and he observed a decrease of the cavity Q-factor, caused by the absorption of photons by the proton spins.

The only real surprise in this work was the thermalisation time, which turned out

much shorter than expected: seconds rather than hours. Purcell invested a significant amount of time over the following years in trying to understand relaxation processes in NMR. He quickly realised that spontaneous emission was a negligible factor since the Einstein A coefficient associated with the proton magnetic moment in free space is of the order of 10^{21} s. However, building on his experience with RF/microwave resonating circuits and two-level atoms, which he acquired during the war while working on radars, he also came to realise that this spontaneous lifetime would be reduced inside the resonant cavity, because more modes are available to decay into. The enhancement factor is (30)

$$f = \frac{3Q\lambda^3}{4\pi^2 V}, \tag{1.19}$$

where Q and V are the cavity Q-factor and volume. f has since become known as the *Purcell factor*. In the cavity used in the 1945 experiment, this enhancement is negligible. But Purcell suggested that by using micrometer-scale metallic particles mixed with the nuclear magnetic medium, the spontaneous lifetime could be decreased down to a few minutes, and therefore observed.

As Table 1.1 shows, Purcell's paper had very little impact for almost four decades. In fact it was never published except as an abstract in the Proceedings of the April 1946 meeting of the American Physical Society (31), where it was eclipsed by Purcell and Bloch's papers on the discovery of NMR. It is very revealing that only in 1981 was it pointed out by Kleppner (32) that the same mechanism would lead to the inhibition of spontaneous emission.

Year	Citations	Notes
1946-1950	2	(all by the initial co-authors)
1951-1960	6	(mainly by the initial co-authors)
1961-1970	10	1970: Drexhage experiments
1971-1980	4	
1981-1990	78	Kleppner's paper published in 1981
1991-2000	445	
2001-2010	1420	
2011-...	367	... and counting

Table 1.1: Citations history of Purcell's original 1946 abstract (30) in the last 65 years. Source: Google Scholars.

Drexhage's monomolecular layers

Spontaneous emission lifetime modification was actually first observed simply by looking at molecules close to a mirror, in a series a beautiful experiments performed by Drexhage in the late 1960s (see (33) for a comprehensive review, or Drexhage's neat *SciAm* article for a more popular version (34)). Although he did not have a resonator per se, the

presence of the mirror nonetheless resulted in constructive or destructive interferences in the dipole emission pattern, which also changed the spontaneous lifetime.

In order to position quantum emitters at a small and well-controlled distance from the mirror, Drexhage used fatty acid molecules floating at the surface of a water tank. He would then dip a glass plate (on which a mirror had been previously deposited) vertically into the tank; upon withdrawing the slide a uniform layer of molecules would stick to it via their hydrophobic end. By repeating the dipping process, up to several hundreds, optically flat molecular layers of identical thickness could be built up. A final layer of dye molecules was deposited to probe the electromagnetic field. Having measured the thickness of a single monomolecular layer (2.64 nm) by observing Fabry-Perot interferences between the mirror and the air-molecule interface, Drexhage could now investigate the field from a few nanometres up to a couple of microns from the mirror, with an unprecedented resolution of about 5 nm (the thickness of molecular bilayer).

For example, by building a stair-like structure of monomolecular layers, he was able to probe directly the standing wave pattern of a laser reflected off the mirror at normal incidence, as well as the evanescent wave that arises from total internal reflection. He observed the change in the emission pattern of dye molecules as they approach the mirror, for both electric and magnetic dipole transitions. By adding a second dye whose absorption frequency overlapped with the emission of the first, he demonstrated radiationless energy transfer from one to the other and thus established the existence of the near field around an excited dipole. Finally, he measured the dyes radiative lifetime (around a millisecond) as a function of distance from the mirror and showed that it varied in accordance with the theory. The quantum yield, which normally nearly impossible to access, can also be determined straightforwardly through this measurement (35).

Beautiful as they are, these pioneering experiments do not quite mark the birth of cavity QED. The lack of a proper cavity limited the lifetime modification to a few tens of per cent, and there was no perspective for coherent manipulation of the dye molecules or of the emitted radiation. In fact we had to wait another full decade for the first cavity experiments to become possible.

The birth of cavity QED

For another decade after Drexhage's work, throughout the 1970s, very little progress was made. Cavity QED only kicked off in the early-mid 1980s, with progress in laser and microwave technology. The first experiments were carried out with Rydberg atoms in the microwave regime, chiefly because it is much easier to achieve a small mode volume, comparable to the wavelength, when the latter is measured in centimetres rather than micrometres. Thus Haroche (36) was the first to observe Purcell enhancement in 1983, using a beam of low-lying Na Rydberg atoms ($23S \rightarrow 22P$ transition at 340 GHz) and a niobium superconducting cavity. The atoms are prepared in the $23S$ state prior to entering

the cavity; if the latter is off-resonant, only $23S$ atoms are detected after the transit. When the cavity is tuned in resonance, $22P$ atoms are detected, showing that the radiative decay $23S \rightarrow 22P$ has been enhanced by the cavity. A similar apparatus using cesium atoms and a planar instead of a concave cavity allowed Kleppner to demonstrate inhibition of spontaneous emission in 1985, with a Purcell factor of about 20 (37). Inhibition was actually first observed by Gabrielse a year earlier, in a rather different setting. He showed that he could suppress the microwave (164 GHz) cyclotron radiation of a single electron in a Penning trap, whose electrodes also form a moderate-Q microwave cavity (38). A maser with a single Rb Rydberg atom acting as the gain medium was demonstrated shortly afterwards by Walther (39), who was also the first to observe the collapse and revival of Rabi oscillations in 1987 (40).

The larger mode volumes and faster spontaneous atomic decay associated with optical transitions made experiments in this regime much more challenging, and they lagged behind for a few years. Suppression of spontaneous emission in the near-infrared was demonstrated by Haroche in 1987 (41) on the caesium $5D_{5/2} \rightarrow 6P_{3/2}$ line at 3.49 μm, in a planar geometry. Then Kimble pioneered the field throughout the 1990s, starting with the observation of one-atom Rabi splitting in 1992 (42), using the caesium $D2$ line at 852 nm. His cavity had a relatively large mode volume (1 mm long, 50 μm waist), although most bulk optics cavities would not get much smaller afterwards. Kimble also demonstrated that a single-atom cavity could be used as a controlled phase gate for quantum optics (43).

Trapping atoms inside the cavity by means of the dipole force (also known as a far off-resonant trap, or FORT) was the next big step (44), followed by cooling of the atoms to the motional ground state (45). Lasing with a single trapped caesium atom was achieved in the same breath in 2003 (46). At the same time, improvements in trapping techniques made it possible to load atoms deterministically into the cavity, for example with a standing-wave conveyor belt (47).

With the new millennium, cavity QED entered a new age: that of coherent control. Kuhn (a former student of Kimble) deterministically produced photons with a predetermined temporal profile by adiabatic passage (26), while Rempe (also from the Kimble school) went further and managed to entangle atoms sitting in two distant cavities. Last but not least are Haroche's latest experiments, in which he send Rydberg atoms flying through a superconducting microwave cavity to add (if the atom was excited) or remove (if in the ground state) a photon from the cavity (1). One can also take the opposite approach and use the cavity field to control the internal state of several atoms (48).

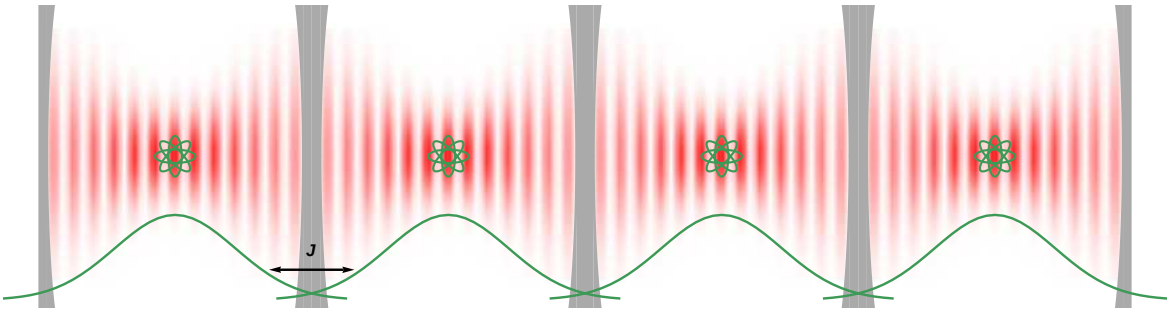

Figure 1.4: A theorist's view of coupled cavities. The coupling J originates from the overlap between the cavity modes (green lines), and an effective photon repulsion U can be generated from one or other variant of the photon blockade effect.

1.4 Coupled-cavity QED

By the mid-2000s, cavity QED had come of age: all the effects discussed in the previous section had been demonstrated and verified beyond the proof-of-principle stage. At the same time much progress had been achieved toward the integration of optical cavities on a microscopic scale with the development of pillar semiconductor microcavities with quantum dots (49), photonic crystal cavities and whispering-gallery resonators. Many-body quantum physics with Josephson junction and, more recently, in optical lattices have also witnessed dramatic developments.

It is in this context that the concept of coupled-cavity QED arose in a seminal 2006 paper by Hartmann and Plenio (50). They proposed to build what is essentially an optical lattice of interacting photons with an on-site non-linear interaction mediated by atoms in cavities, as depicted in Figure 1.4. An array of resonators would be coupled to atomic ensembles; the atoms are four-level N systems where an effective photon-photon non-linear interaction is generated via an EIT-based scheme. The concept is formally similar the Bose-Hubbard model, which has been a paradigm of many-body physics for several decades.

Before we survey some of the many schemes that have been proposed since that paper, let us briefly review the main features of the Bose-Hubbard (BH) model.

1.4.1 The Bose-Hubbard model

The BH model is a very general description of bosons interacting on a lattice. It was proposed in 1989 as an extension of the Hubbard model that deals with fermions (electrons) in semiconductors. The citation history of the original BH paper (51) shows a marked increase in the rate of citations in the mid-2000s (from less than 50 to over 200 per year) that coincides primarily with the advent of cold atoms experiments on optical lattices, but is also related to the interest in coupled-cavity QED.

The BH model considers a periodic potential in which each site i is populated with bosons (annihilation/creation operator $\hat{\Phi}_i, \hat{\Phi}_i^\dagger$, number operator $\hat{N}_i = \hat{\Phi}_i^\dagger \hat{\Phi}_i$, with the

standard bosonic commutation relation, $[\hat{\Phi}_i, \hat{\Phi}_j^\dagger] = \delta_{ij}$). Bosons experience on-site repulsive potential U while being able to hop between neighbouring sites i and $i \pm 1$ at a rate J_i. The on-site energy of a single particle is given in term of a chemical potential μ. The BH Hamiltonian then reads

$$H = \underbrace{-\sum_i \mu \hat{N}_i}_{\text{on-site energy}} + \underbrace{\frac{1}{2} V \sum_i \hat{N}_i(\hat{N}_i - 1)}_{\text{repulsive interaction}} - \underbrace{\frac{1}{2} \sum_i J_i(\hat{\Phi}_i^\dagger \hat{\Phi}_{i+1} + h.c.)}_{\text{boson hopping}}. \tag{1.20}$$

The on-site interaction is proportional to $\hat{N}_i^2$, and is therefore akin to a Kerr non-linearity in optics.

In the absence of disorder, the physics governed by such a Hamiltonian can be very intuitively understood in the two limits $U/J \gg 1$ and $U/J \ll 1$.

- If $U \gg J$, the repulsive interaction U dominates and effectively prevents excitations from tunnelling from one site to the next; they are strongly localised at each lattice site. This quantum phase is the counterpart of a crystal in condensed matter, and is called a *Mott insulator*.

- If $U \ll J$, the tunnelling dominates and the bosons, free to circulate, are delocalised across the entire network, forming a *superfluid phase*.

Whether the system is in one or the other of these phases can be experimentally determined by observing the fluctuations in the photon numbers inside each cavity, which would be suppressed in the Mott insulator phase. Such a quantum phase transition was first observed with a BEC held in an optical lattice (52), in which the tunnelling rate J can be tuned by changing the depth of the lattice potential, i.e. the trapping laser intensity. However optical lattices lack the ability to tune the coupling between two sites independently of the others, and the sites are too close to each other to enable single-site manipulation. In the presence of disorder (e.g. random variations of U or J between sites, the existence of an exotic "glassy" phase has been predicted between the Mott insulator and superfluid phases (53).

1.4.2 Theoretical proposals

The first important step in building a coupled-cavity QED system is to generate a strong enough effective photon-photon interaction. The simple Jaynes-Cummings photon blockade effect is a candidate, but it suffers from decoherence via spontaneous emission from the excited state. This can be alleviated to some extent by detuning the cavity from the atom to work in the dispersive regime, but the non-linearity is then smaller, and decreases even further with increasing number of atoms (54). An alternative is to use 4-level EIT-like schemes, in which the excited state population is small enough that spontaneous

emission can be neglected. Such a scheme was initially proposed by Imamoğlu (55; 56). A slightly simpler concept was devised by Hartmann (54) using only a three-level atom, in which the Stark shift induced by the cavity field on one transition generates the required non-linearity. Many variants of these basic schemes have been studied (57; 58). In general, the strength of the non-linearity can be increased by adding more atoms in the cavity.

The first coupled-cavity QED proposals were formulated in two back-to-back articles in the same 2006 *Nature* issue, by Hartmann (50) and Greentree (59). The first one was based on a traditional cavity design and 4-level atom EIT non-linearity, where the coupling rate J is proportional to the intensity of an external driving laser. The second one considered a two-dimensional network of diamond photonic crystal cavities with NV centres as two-level atoms, with the ability to tune the interaction strength U by detuning (Stark-shifting) the emitters from the cavity. In both cases, the bosonic excitations were not purely photons but mixed atomic-photonic excitations referred to as polaritons. It was subsequently demonstrated that the polaritons could, in the appropriate regime, be almost pure photons and therefore a *photonic Mott insulator* was possible (60).

A three level atom in a cavity can be modelled as a spin, and therefore a system of coupled cavities can implement an effective spin chain (61; 62). The high degree of control over the parameters make it an attractive system for exploring spin dynamics; for example Illuminati showed that by tailoring the coupling rates along a very long spin chain, robust entanglement between the two ends could be created and teleported (63). Specifically, he considered a so-called $\lambda - \mu$ chain in which the couplings are arranged as $\lambda - \mu - 1 - 1 - 1 - ... - 1 - 1 - 1 - \mu - \lambda$, with $\lambda < 1 < \mu$. In general, thanks to their high degree of control and flexibility, arrays of coupled cavities have been primed as a good candidate for quantum emulators and even as a cluster state generator for quantum computing. For a comprehensive review, we refer the readers to (64).

In light of recent experimental efforts, including ours, it became clear that realistic coupled-cavity QED implementations would have to accommodate relatively large losses. Although still limited in scope, there has been a renewed interest in dissipative systems indicating that non-trivial quantum physics may be observable even in the presence of strong dissipation (65; 66; 67; 68; 69).

1.5 Practical implementation

1.5.1 CCM's efforts toward integration

The work we undertook in this thesis is, to our knowledge, the first attempt to implement the coupled-cavity QED proposals outlined in the previous section. It drew largely on previous (and still ongoing) work at the Centre for Cold Matter (CCM) on microcavities

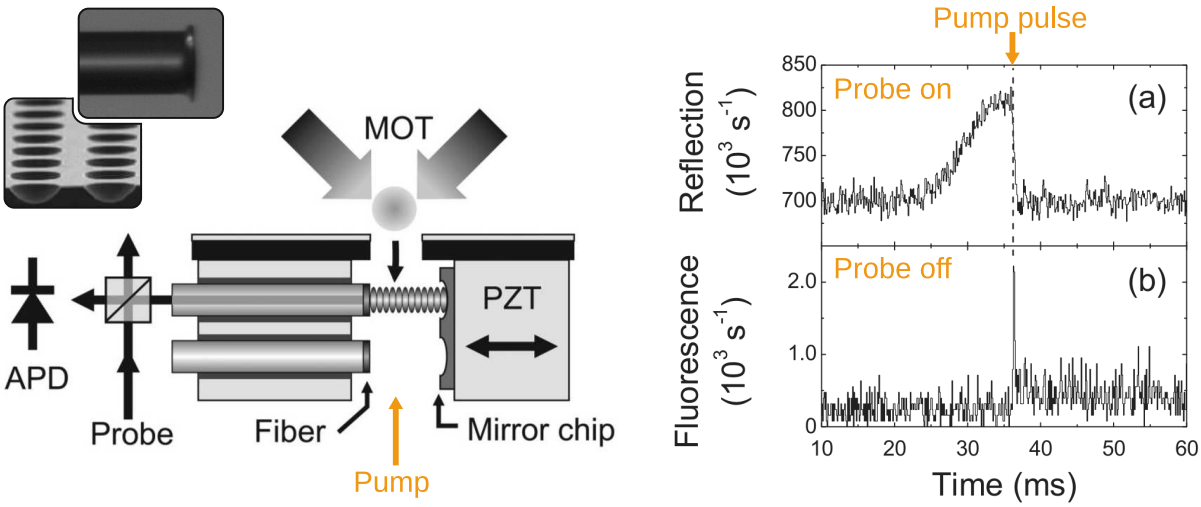

Figure 1.5: The microcavity QED experiment at CCM. Left, the experimental set-up. Micrographs of the fibred mirror and of the spherical micromirrors array are shown in inset. Right: Experimental results. Ground state Rb atoms are dropped from a MOT and fall freely through the cavity mode; the passage of a single atom disturbs the cavity, as seen in the change in the reflection of a probe beam in the first half of trace (a). The maximum mean atom number is $\langle n \rangle \simeq 1$. At the time indicated by the vertical dashed line, the atoms are excited by the pump laser shown in the diagram (left) into the excited state; a sharp burst of cavity-enhanced photon is then observed through the fibre (b). About 50% of the photons are emitted into the cavity mode, indicating a cooperativity $C \sim 1$. Graphics adapted from (71).

and waveguide atom chips, which we review first.

Fibre-coupled microcavities

Most of the cavity QED experiments we discussed in Section 1.3.2 were based on bulk optics. Very high finesses were achieved at the expense of a relatively large mode volume and very limited prospect for scalability. It is with this latter challenge in mind that in the mid-2000s the CCM group started developing arrays of spherical micromirrors etched on a silicon substrate and coated with a high-reflectivity multilayer dielectric mirror (70). A cavity is formed between this mirror and another planar mirror glued on the tip of an optical fibre. Finesses of up to 5000 have been measured. The spherical mirrors have a typical radius of curvature in the range $\sim 100 - 300\,\mu$m, which allows the cavity mode to be perfectly matched to the fibre. Such a cavity has a mode volume at least an order of magnitude smaller than bulk optics designs, with a waist of about $5\,\mu$m and a length of order $100\,\mu$m.

These cavities can be used as a fast, efficient single-atom detector (72) and as a single photon source (71), as depicted in Figure 1.5. Cold rubidium atoms from a MOT are dropped in free fall into the cavity, which is probed with a weak resonant laser. In this case we operate in the fast cavity regime $\{\kappa = 13 \gg g_{\rm ac} = 0.6 \gg \gamma = 0.02\}$GHz, giving a theoretical cooperativity $C \simeq 0.6$. If the atoms are in their ground state, they act

at resonance as scatterers/absorbers; we observe this as a reduction in the cavity fringe contrast in the first half of the trace in Fig. 1.5(a). When the atoms are excited by a pump laser (Fig. 1.5, left), the reflection signal vanishes because the recoil momentum pushes the atoms out of the cavity. However if we repeat the experiment without the probe beam, we detect a sharp burst of photons coming out of the fibre, as shown in Fig. 1.5(b). These are spontaneously emitted by the atoms with some enhancement from the Purcell effect: we can calculate that roughly 50% are coupled into the cavity mode and detected out of the fibre; thus the cooperativity is about 1, as expected. Normal-mode splitting has also been observed.

The micromirror array may contain several hundred mirrors. Fibres can also be integrated in large, commercially available arrays of v-groove, or into a waveguide chip; thus is should be possible to operate a large number of cavities simultaneously.

Array of integrated atom–photon junctions

In a second experiment, twelve waveguides were brought into close proximity in the middle of a chip (10 μm apart), and a 16 μm-wide trench was cut perpendicular to the waveguides (73). Atoms were loaded in the trench, and each waveguide could be used to probe the atomic cloud at a different position. The atoms were used to measure the light polarisation, and the waveguides could resolve the Gaussian profile of the atom cloud. In the future it may be possible to do experiments with a BEC in which a wide range of excitations could be induced via the waveguides. A periodic two-colour potential, with light at two different frequencies in alternate waveguides, could be used as an atom mirror.

Based on the experience accumulated during these two project, we envision an array of Fabry-Pérot microcavities coupled not to fibres, but to waveguides on a chip. But before we describe it, let's have a look at possible alternative.

1.5.2 Possible implementations of coupled-cavity QED

Photonic crystals

Photonic crystals (PC) (74) are a generalisation of multilayer dielectric (Bragg) mirrors, which are just 1D PCs. They rely on Bragg reflection on periodic dielectric structures to confine or guide light in a way that is very similar to electrons in semiconductors; hence they are also often called photonic bandgap materials. In their most common form, PCs are thin slabs (~ 100 nm) of high-index dielectric patterned with a regular array of holes. A missing row of holes defines a waveguide and a few missing holes a cavity, so that one can conceive of an array of PC cavities close enough to be coupled via their evanescent field. This is the setting for most of Greentree's theoretical work.

Coupled photonic crystal cavities have already been experimentally demonstrated. In 2008, NTT Labs in Japan fabricated a waveguide along which they defined identical,

closely spaced cavities by tweaking the spacing between holes (75). They showed normal mode splitting between up to 200 co-resonant cavities with Q-factors in excess of a million and a mode volume of $1.6(\lambda/n)^3$. The theoretical Purcell factor is then several tens of thousands.

More recently, the Noda group also in Japan (pioneers of PCs) built a slightly different system. They defined a long waveguide, closed by reflectors and evanescently coupled to a cavity at each end. The length of the waveguide is such that it is exactly off-resonant with the cavities, which are too far away to be directly coupled. But the waveguide nonetheless induces a coupling between the cavities, and because it is non-resonant it is only weakly populated — such a configuration should be able to accommodate some propagation losses. They observed both normal mode splitting and "photonic Rabi oscillations" between the two cavities at a rate of 18 GHz (76).

Photonic crystals are thus a very promising platform for integrating multiple cavities. However, there is as yet no suitable quantum emitter. Quantum dots (QDs) are obvious candidates, and they have been successfully integrated to single cavities (77), but for our purpose we need to position them accurately within each cavity. Moreover the exciton transition needs to be degenerate across many QDs. Currently the technology is lacking on both counts. Moreover, the coupling rate between cavities is set by the geometry and is difficult to tune. Nonetheless, given the continued interest in quantum dots and photonic crystals, the situation may change in the future. In a promising recent experiment from the Solomon group (78), a QD was coupled to the symmetric mode of a "photonic molecule" made of two coupled PC cavities, while the antisymmetric mode was pumped to induce a cavity-enhanced Stark shift large enough to tune the QD in and out of resonance with the symmetric mode. This could be a mechanism to tune the QD-cavity coupling, and consequently the photon-photon interaction strength U.

Toroidal microcavities

A very different kind of microcavity relying on whispering-gallery modes can be made, either in a toroidal geometry (79) or in silica sphere fabricated by melting the tip of an optical fibre (80). Although the mode volume is larger than in photonic crystals, this is compensated by a very high quality factor thanks to the extreme surface smoothness and high material purity that can be achieved.

Such resonators are coupled via a tapered nanofibre, and it is possible to couple several of them to a single fibre (81). Quantum dots can be deposited on the surface and coupled to the evanescent field, or a cold atomic ensemble can be brought in close proximity; atoms could even be trapped in the same way as with nanofibres (82). The relatively large distance between cavities allows individual addressing of the emitters.

Superconducting microwave resonators

In the microwave regime, superconducting stripline resonators offer both large Purcell factors (because of the long wavelength — see Eq. (1.19)) and very high Q-factors. Experiments have been performed showing strong coupling between ensembles of diamonds NV centres (83) with a stripline resonator. One can imagine that several of these systems could be coupled together. It is also possible to use "artificial Q-bits", such as Josephson junction and Cooper-pair boxes (84).

1.5.3 Array of waveguide-coupled Fabry-Pérot microcavties

Our own original design was evolved in collaboration with Michael Hartmann and Martin Plenio (now respectively at the Technische Universität München and at the University of Ulm) as part of the European project HIP (Hybrid Information Processing). According to the programme's statement of purpose (85):

> Scaling quantum information processors beyond the present small-scale devices is challenging as communication between parts of the processor, single site addressability and scaling are difficult to reconcile.
>
> HIP addresses these issues with the experimental realization of elementary hybrid atom-photon devices, and the theoretical development of schemes for their integration on platforms capable of being miniaturized and scaled up in functional networks. The main experimental platform on which this goal will be pursued are atom chips structures on which optical micro-cavities will be integrated. These devices will then be connected with optical fibres to form a network.

Device description

Along these lines, we devised the coupled-cavity QED system sketched in Figure 1.6. The emitters sit in microcavities formed between a spherical mirror R_{c} and a plane mirror R_{cw} deposited on the end facet of a waveguide chip . The microcavities are open in the transverse direction, giving access to lasers to trap and manipulate atoms at the position of maximum interaction with the cavity modes. The concave mirrors are the same as those currently used in the CCM microcavity experiment (86), and are fabricated on a silicon substrate by wet isotropic etching (70) before being reflection-coated. Large arrays of regularly spaced mirrors can be made. With additional polishing we expect the losses due to surface roughness to be of order $A_{\mathrm{c}} \approx 10^{-5}$ (87), while the plane mirror (subscript cw) can have losses on the order of $A_{\mathrm{cw}} \approx 5 \times 10^{-6}$ if the waveguide facet is super-polished and a dielectric mirror formed by ion-assisted deposition is used.

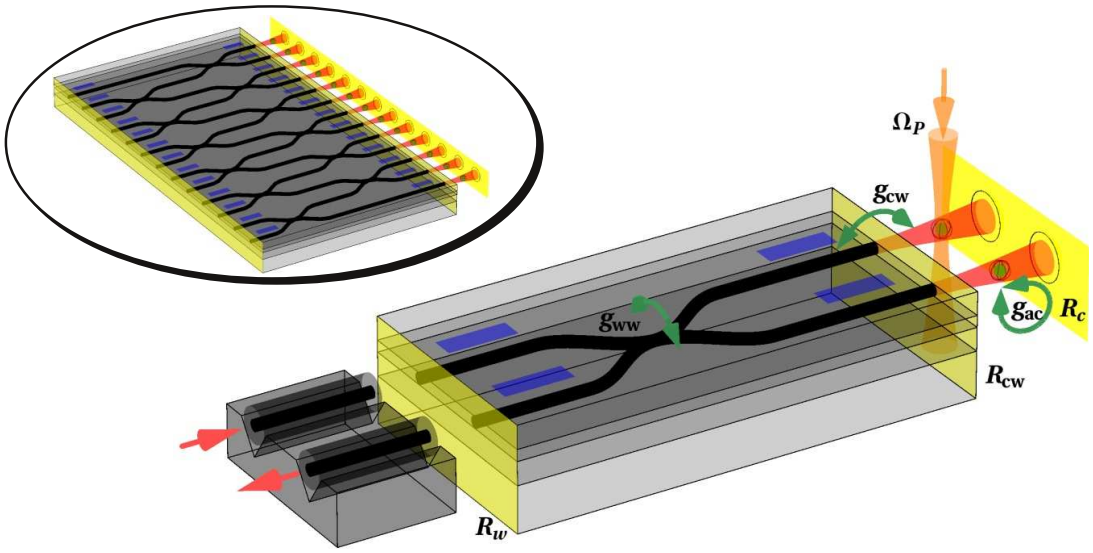

Figure 1.6: Realistic, although not to scale, representation of the proposed coupled cavity array: a basic 2×2 unit showing mirror reflectivities, coupling rates (atom-cavity $g_{\rm ac}$, cavity-waveguide $g_{\rm cw}$, waveguide-waveguide $g_{\rm ww}$), and side pumping Ω_P. A V-groove array of fibres is used to connect the device. The concept is highly scalable (inset).

The reflectivity of each mirror is given by $R_i = 1 - (T_i + A_i)$, where T represents the power transmission and A the loss. Assuming that $T_i, A_i \ll 1$, the cavity field amplitude decays at a rate (see also Eq. (3.11))

$$\kappa_{\rm c} = \frac{c\xi_{\rm c}}{2L_{\rm c}} \quad \text{with} \quad \xi_{\rm c} \approx \frac{T_{\rm c} + A_{\rm c} + T_{\rm cw} + A_{\rm cw}}{2}, \tag{1.21}$$

which could be made as small as $\kappa_{\rm c} \approx 2\pi \times 0.01\,\text{GHz}$ for a cavity of length $100\,\mu$m. This will give enough space for external optical access to the atoms. For the latter, we consider the D_2 line of rubidium-87 at 780 nm, whose amplitude decay rate is $\gamma = 2\pi \times 3\,\text{MHz}$. The cavity waist could be as small as $2\,\mu$m, and therefore the atom-photon coupling rate or Rabi frequency is of order $g_{\rm ac} = 2\pi \times (0.1 - 1)\,\text{GHz}$. Thus the strong coupling regime $g_{\rm ac} \gg \gamma, \kappa_{\rm c}$ is in principle achievable. These questions are discussed more fully in Chapter 3.

We then define waveguides on the chip, aligned with the microcavities. These waveguides can be evanescently coupled to their neighbours, and in order to confine photons within the system we add a mirror $R_{\rm w}$ on the second, external end facet of the waveguide chip. Thus the waveguides themselves form a second set of (coupled) resonators, with a hoping rate $g_{\rm ww}$ (see §4.2.1). The coupling between the microcavity and waveguide resonators is described by another coupling rate $g_{\rm cw}$ (see §4.1.2). For nearest-neighbour coupling, we want to operate in a regime where

$$R_{\rm c} \gg R_{\rm w} > R_{\rm cw} \quad \text{and} \quad g_{\rm cw} \gg g_{\rm ww}. \tag{1.22}$$

Thus a photon emitted in a first microcavity will hop to the waveguide resonator rather than get lost in the silicon substrate, then tunnel to the next waveguide – and into the adjacent microcavity – instead of immediately leaking out of the system.

Finally, the waveguides can be interfaced with an array of optical fibres (arrays of up to 64 fibres are commercially available (88)), so that photons can be efficiently collected (much more so than with photonic crystals for example). Pump light can also be injected in the system through this interface.

The device does not need to be entirely passive. In fact, it will be necessary to tune the coupling rates in order to observe quantum phase transitions. With our design this is very easy to achieve, provided two phase shifters can be integrated in each waveguide, one on each side of the couplers (blue in Fig. 1.6). This allows us to tune g_{ww}, as demonstrated in §4.2.3. The inset in Fig. 1.6 highlights the intrinsic scalability of the device.

Theoretical overview

A detailed classical optics analysis of the device is presented in Chapter 4. Its quantum dynamics has been studied in collaboration with Michael Hartmann and Martin Plenio, published in the *New Journal of Physics* (89). We refer the reader there for the details; here we only summarise the main conclusions.

1. The combined microcavity-waveguide resonator system can be regarded as a basic unit and the array of cavities can be described by coupling between single modes of these units even though they have a complex spectrum (which is fully derived in Chapter 4.1), provided that the coupling g_{ww} is much smaller than the waveguide free spectral range. Thus, despite the additional waveguide resonator stage, this structure can be described in exactly the same way as a simple 1D chain.

2. An elementary two-cavity Jaynes-Cummings-Bose-Hubbard model is considered, where a single two-level atom resides in each of the two cavities. The driven steady-state spectrum of such a system exhibits four peaks corresponding to four eigenmodes. When properly tuned, the second-order correlation function reveals strong antibunching, proving that the system can generate non-classical light. We also show that photons in the two cavities become entangled. These effects are observable with standard mirror reflectivities: $R_{\mathrm{c}} = 99.9\%$, $R_{\mathrm{cw}} = 98.0\%$, $R_{\mathrm{w}} = 99.8\%$.

3. We then consider an effective spin Hamiltonian with two cavities containing a single three-level atom and study the spin dynamics. With the previous parameters, the system is over-damped. However this can be rectified by increasing the reflectivities to $R_{\mathrm{c}} = 99.99\%$, $R_{\mathrm{cw}} = 98.0\%$, $R_{\mathrm{w}} = 99.9\%$, which is not unreasonable. Coherent spin oscillations between the cavities can then be observed, with a coherence time limited by dissipation in the waveguide resonator.

1.5.4 Thesis outline

In this introductory chapter, we have explained the deep motivation behind cavity QED, reviewed its historical development as well as the latest advances, and motivated "coupled-cavity QED" as its logical extension with potentially far-reaching application for quantum information science. We have proposed an original and realistic implementation and outlined its theoretical capabilities.

The following three chapters describe in details our efforts towards the practical realisation of the proposed device. Chapter 2 covers the fabrication of the required waveguides, couplers and phase shifters. Chapter 3 considers the waveguide resonators and the microcavities separately; they are brought together in Chapter 4. In Chapter 5 we present our work on Bragg gratings which we had hoped to integrate to the waveguides.

The limitations of the current systems and prospects for the future are laid out in the Conclusion.

Chapter 2

Waveguide chips: theory, fabrication and characterisation

2.1 Guided optics: An introduction

All the waveguides used in this work are, conceptually, optical fibres on a chip, with core diameters around 5 μm and an index contrast of the order of 10^{-3}. Optical fibres exploiting the low intrinsic loss of fused silica at wavelengths close to 1550 nm have become ubiquitous in the modern world, enabling ultra-broadband, long-distance communications. Charles Kao, who spent decades pushing forward what started as an outsider technology, was rewarded in 2009 by a physics Nobel Prize for "groundbreaking achievements concerning the transmission of light in fibres for optical communication". It was a long road: in the 1950s guided millimetre waves were in favour for telecommunications, and none of the major companies (with the exception of the British Post Office) would invest in optical fibres, whose propagation losses were initially in the hundreds of dB per kilometre.

From 1966, when the potential for optical communication was established by Kao (his estimate of the achievable bandwidth has now been exceeded by 5 orders of magnitudes, and he envisioned that losses could be brought down to 20 dB/km), progress was rapid. On the theoretical front, Snitzer solved Maxwell's equation for a cylindrical dielectric waveguide in 1961 (90); a decade later in 1971 Gloge devised the weak guiding approximation (91) while Goell developed the first numerical tools for non-cylindrical geometries in 1968 (92). The first viable fibre communication link was set up in 1965, and commercial links appeared a decade later. In 1977 NTT started using the "third transparency window", at 1550 nm, where the propagation loss is only 0.2 dB/km. Erbium fibre amplifiers were developed in the early 1990s, enabling transoceanic links.

Planar optical waveguide chips developed on a parallel road, initially as an integral part of semiconductor laser technology (the gain region in a semiconductor laser often acts as a light guide). The first solid-state laser was demonstrated in 1960, winning Townes

the 1964 Nobel Prize; and the first semiconductor laser was build in 1962 (93), eventually leading again to a Nobel prize in 2000, awarded to Alferov and Kroemer "for developing semiconductor heterostructures used in high-speed- and opto-electronics". Thus the development of planar optical circuits was also closely linked to the telecom industry, and the silica-on-silicon technology, largely derived from optical fibre fabrication, provided many of the elements needed for switching (using the thermo-optic effect) and wavelength division multiplexing with arrayed waveguide gratings (94). Where a high switching rate was required, electroactive devices based on lithium niobate provided GHz bandwidth (95).

Total internal reflection

After this brief historical introduction, we turn our attention to the theoretical aspects of wave guiding. The basic underlying mechanism is total internal reflection (TIR), which was first studied back in 1842 when Colladon observed that sunlight was guided inside a jet of water. Snell and Descartes in the seventeenth century, and before them the Persian Ibn Sahl in the tenth century (96), had discovered that the deflection (refraction) of a light ray at the boundary between two dielectrics of refractive index n_1 and n_2 is governed by the law $n_1 \sin i_1 = n_2 \sin i_2$ ($i_{1,2}$ are the angles of the incident and transmitted beams relative to the normal). Thus a ray crossing an interface from the optically denser medium ($n_1 > n_2$) will be deflected closer to the surface, and since i_2 cannot be larger than $\pi/2$, there is a critical incident angle $i_\mathrm{c} = \arcsin(n_1/n_2)$ beyond which all of the incident beam will be reflected (Fig. 2.1(a)). The electromagnetic field in the second medium does not vanish completely, but forms an *evanescent wave* that decays exponentially away from the interface (its k-vector is imaginary). TIR is treated extensively in many electromagnetism textbooks (97), and is an extremely convenient way of analysing massively multimode waveguides in particular. If instead of a single plane interface we have a closed surface delimiting a channel of an optically denser medium, it will act as a light guide. For large multimode waveguide, such a ray analysis is sufficient. For smaller, single-mode waveguide such as those we are interested in, we require a more detailed wave analysis based on Maxwell's equations.

In the following, we will briefly review the standard derivation of the guided modes of an optical fibre.

Cylindrical waveguide

Consider a cylindrical waveguide of radius a (Fig. 2.1(b)) in cylindrical coordinates $\{r, \theta, z\}$, with core and cladding indices n_co and n_cl. For TIR guiding we require $n_\mathrm{co} > n_\mathrm{cl}$. From Maxwell's curl equations it is possible to express E_r, E_θ, H_r and H_θ as functions of E_z and H_z only (98). Thus we only need to solve the wave equation $(\nabla^2 + k(\mathrm{r})^2)\{\mathbf{E}, \mathbf{H}\} = 0$ for these two components, which is relatively simple because the unit vector $\mathbf{z}$ is constant.

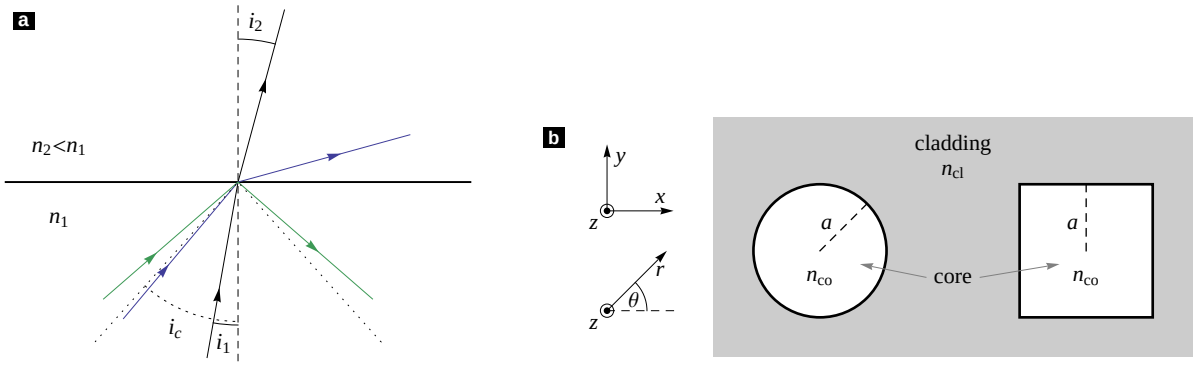

Figure 2.1: (a) Refraction and total internal reflection to an optically less dense medium. Rays are deflected away from the normal, so for $i_1 > i_c$ reflection will occur (the critical angle i_c is represented by the dotted line). **(b)** Cartesian and cylindrical coordinate system. Cylindrical and square waveguides.

The problem is fully separable, so we can write $E_z(\mathbf{r}, t) = R(r)\Theta(\theta)e^{i(\omega t - \beta z)}$, and similarly for H_z. The wave equation then reduces to a Bessel differential equation with solutions of the form

$$\Theta(\theta) = \sin(n\theta + \phi) \tag{2.1}$$

$$R(r) = \begin{cases} AJ_n(ur/a) & 0 < r < a \quad (n \in \mathbb{N}) \\ BK_n(wr/a) & r > a \end{cases} \tag{2.2}$$

with $u = k_0 a\sqrt{n_{\text{co}}^2 - n_{\text{eff}}^2}$, $w = k_0 a\sqrt{n_{\text{eff}}^2 - n_{\text{cl}}^2}$ and $n_{\text{eff}} = \beta/k_0$ such that $n_{\text{co}} > n_{\text{eff}} > n_{\text{cl}}$. In the cladding these solutions correspond to evanescent waves, which behave asymptotically as $K_n(r \gg 1) \sim e^{-r}/r$. n_{eff} is the modal index and c/n_{eff} the phase velocity of the guided mode. The integer parameter n is the number of azimuthal field nodes.

At this stage there are four unknown coefficients, for E_z and H_z in the core and in the cladding. To determine these, we first have to express the two other tangential components E_θ and H_θ as functions of E_z and H_z. All of these must be continuous at the core-cladding interface $r = a$, which results in a set of four continuity equations whose determinant has to vanish to yield non-trivial solutions. This requirement is written as a characteristic equation

$$\left(\frac{J_n'}{uJ_n} + \frac{K_n'}{wK_n}\right)\left(\frac{n_1^2}{n_2^2}\frac{J_n'}{uJ_n} + \frac{K_n'}{wK_n}\right) = n^2\left(\frac{1}{u^2} + \frac{1}{w^2}\right)\left(\frac{n_1^2}{n_2^2}\frac{1}{u^2} + \frac{1}{w^2}\right) \tag{2.3}$$

which can be solved for n_{eff}. It admits a small number of solutions defining the allowed guided modes. Their number is determined by the *normalised frequency* $v = \sqrt{u^2 + w^2} = k_0 a\sqrt{n_{\text{co}}^2 - n_{\text{cl}}^2}$. For $v < 2.405$, their is only one solution (for $n = 1$): this is the single mode regime. For a given n, solutions with increasing n_{eff} (or, equivalently, v) correspond to 0, 1, 2... field nodes in the radial direction.

For $n = 0$, there are two families of modes with $E_r = E_z = H_\theta = 0$ or $H_r = H_z =$

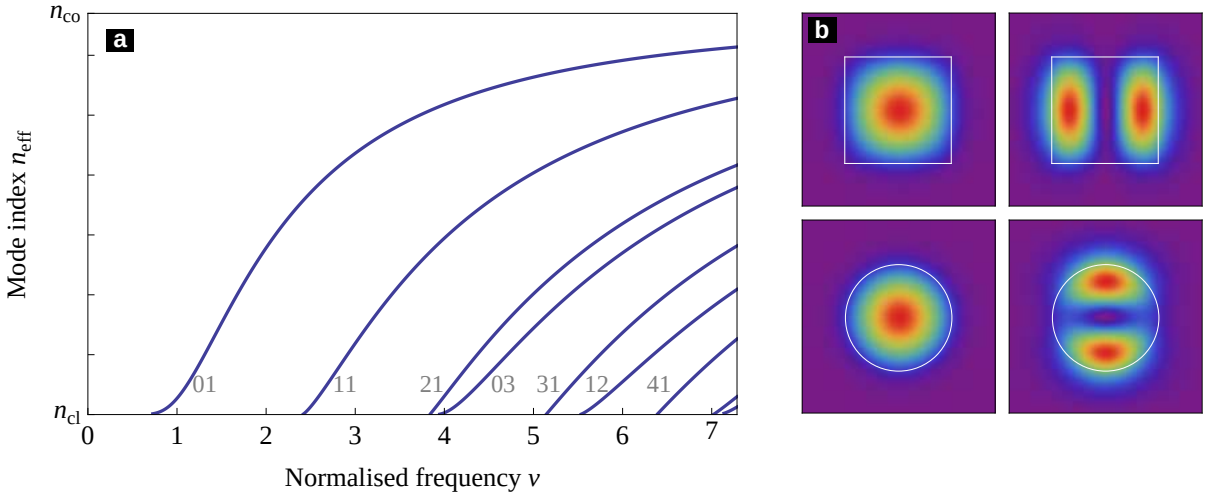

Figure 2.2: Effective index of weakly guiding cylindrical waveguide as a function of normalised frequency v, together with the LP_{mn} mode numbers. Right, selected mode functions of $v = 4.3$ ($a = 5\,\mu$m, $\Delta n = 0.004$) square (top) and cylindrical (bottom) waveguides.

$E_\theta = 0$, known as transverse electric and magnetic modes (TE /TM, **E** or **H** being purely transverse). For $n \geq 1$ the fields are more complex, as all of the components are non-zero. They are called EH or HE modes, depending on the relative weight of their E_z and H_z components. However, in the *weak guiding limit* $n_{\text{co}}^2 - n_{\text{cl}}^2 \simeq 2\Delta n n_{\text{co}} \ll 1$, many of these EH/HE modes become degenerate, and the new modes that arise form their combinations are very well approximated by so-called linearly polarised (LP) modes. They can be conveniently expressed in a cartesian basis, and fall again in two orthogonal families with either $E_x = H_y = 0$ or $E_y = H_x = 0$, while the z-component is very small (hence the denomination).

The modal index in the weak guiding regime is plotted in Figure 2.2 as a function of the normalised frequency, together with a few representative mode functions.

Other waveguide geometries

The derivation thus outlined is quite general, but only works when all the variables can be separated. In general this is only possible for an infinite slab waveguide and cylindrical/ellipsoidal fibres. Most other geometries do not admit analytical solutions, so we must resort to numerical tools. Various methods have been developed over the years. Some are semi-analytical: Goell's expansion of the modes of a rectangular waveguide in terms of the modes of a cylindrical one (92), Fourier expansion in sine waves (99), "modified" Fourier expansion (which adds a mapping of the infinite plane to a unit square to avoid imposing an arbitrary bounding box) (100). Today, with the large memory and computational power that has become widely available, purely numerical finite-difference simulation have become more common. We used MIT's open-source packages, MEEP (time-domain) and

MPB[1] (frequency-domain) (101; 102). The mode profiles in Figure 2.2 have been calculated using MPB.

Optical fibres mode are very well approximated by Gaussian mode. In fact the mode overlap η^2 (which determines how much power can be converted from one mode to the other; see Appendix A for more details) between the fundamental LP mode of a fibre of radius a and a Gaussian mode of waist w_0 can be as high as 99.8%; this optimum is reached when (103)

$$\frac{w_0}{a} = 0.65 + 1.619v^{-3/2} + 2.879v^{-6}. \tag{2.4}$$

Note that even though the Gaussian approximation is extremely good and useful in and around the core, it breaks down further away since the $K_n(r)$ functions scale like e^{-r}/r and not like e^{-r^2}. Therefore it cannot be used to describe evanescent coupling.

2.2 Waveguide fabrication

After this brief introduction to the physics of waveguiding and optical fibres, we now describe the fabrication and characterisation of the waveguides used throughout this thesis. Most of them were made in collaboration with Peter Smith's group at the Optoelectronics Research Centre of the University of Southampton, by direct UV-writing in a silica-on-silicon platform (104). But we also used a femto-second-written chip graciously given to us by the Ultrafast Optics group of the Ultra Optics Centre at the Friedrich Schiller University in Jena, Germany (105), as well as a more traditional deposition-and-etching silica-on-silicon chip (CIP Technologies) borrowed from CCM's Waveguide experiment (106).

2.2.1 Deposition and etching: CIP chip

To this day, the most common fabrication technology for silica-on-silicon waveguides is the deposition-and-etching process. As its names indicates, it consists in depositing layers of appropriate materials by one of many processes (CVD, PECVD, FHD, SolGel, thermal oxidation...) followed by the removal of most of the material to define waveguides, and finally – if needed – the deposition of a cover layer.

The CIP chip was fabricated as follows.

First, a 10 μm thermal oxide layer was grown on the silicon substrate. Silicon has a high affinity for oxygen and therefore will naturally form a thin SiO_2 layer. Under high pressure, high temperature (1000°C, 25 atm) and in an wet, oxygen-rich atmosphere, this silica layer can reach a thickness of 5 to 20 μm. Thermal oxidation is a very common

[1]MEEP and MPB are Linux-only software. Some people have attempted running it on Cygwin. Now they know better.

way of producing the underclad layer; it has the additional advantage of helping releasing strain, thanks to the smooth transition between the silicon and oxide layers.

The second step is to deposit doped, higher-index silica to form the core layer. At CIP as well as in Southampton, this is done by Flame Hydrolysis Deposition (FHD), one of the most common integrated circuits process, in part because it was already a mainstream technology for the production of optical fibre preforms. A silicon tetrachloride ($SiCl_4$) vapour is injected in an oxygen-hydrogen flame to produce "soot", tiny silica particles. Other vapours can be simultaneously injected to incorporate dopants to the soot: $GeCl_4$ and BBr_3 will add germanium and boron, respectively. The concentration of these dopants can be finely tuned to achieve the desired refractive index, with Δn up to a few per cent. The soot is deposited on a substrate located immediately beneath the flame, and is then consolidated at high temperature to form a fully dense silica layer. Several layers with different doping can be deposited consecutively.

After the doped FHD layer, a UV photoresist is deposited and exposed to define waveguides. Most of the FHD layer is then etched away by reactive ion etching (RIE), leaving only waveguide cores. Finally, a second FHD layer, index-matched to the thermal oxide, is deposited to form the upper cladding. Losses as low as 0.05 dB/cm at 1550 nm have been reported (107).

The CIP chip incorporated 14 waveguides, 250 μm apart at the end facets and coming together as close as 10 μm in the centre. Many identical chips were made, and most of them had a trench drilled in the middle, perpendicular to the waveguides; atoms were then brought into the trench to make a compact "array of integrated atom-photon junctions" (73). The chip we used however did not have the trench.

FHD is not the only technique available for making deposition and etching waveguide devices, although it is one of the most common. There are many others (CVD, PECVD, FHD, SolGel...), but they all required expensive and complicated equipment and many different steps. The process is therefore only really suitable for industrial mass production, but not for small-scale applications in a research environment where cheap and fast prototyping is desirable.

Fortunately, there are alternative techniques that offer such advantages: directly written waveguides, by continuous UV or pulsed IR radiation.

2.2.2 UV-written waveguides: ORC chips

It was discovered in the 1970s that the refractive index of germano-silicate glasses was sensitive to UV radiation exposure. The process is still not very well understood due to the complexity of the material: its amorphousness, the presence of many different dopants with different relationships to the surrounding glass, and also the difficulty of performing precise deep UV spectroscopy. Two classes of effects are usually considered. At the

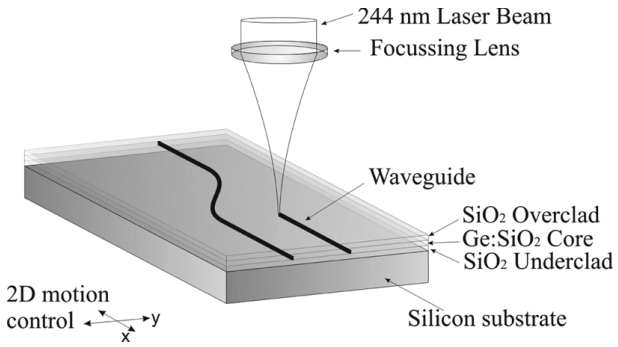

Figure 2.3: Schematics of direct UV-writing of waveguides in a silica-on-silicon chip. The middle Si layer is photosensitive thanks to germanium doping. Drawing courtesy of James Gates/ORC.

microscopic level, the germanium sites in silica have a strong transition near 242 nm, and the energy absorbed from a UV laser is then transferred to other sites nearby where defects are created, leading to an increase of the refractive index. At the macroscopic level, UV exposure compactifies the glass, also increasing the index. The details of both the micro- and macroscopic effects, as well as their relative contributions, are still a subject of research and debate (108).

The first application of direct UV-written was the fabrication of Bragg gratings inside germanium-doped optical fibres in 1989. Silica-on-silicon waveguides were written using this technique in 1994 (109), although similar waveguides had been demonstrated much earlier in polymer substrates. Basic devices such as directional couplers were demonstrated a few years later, and work started at the ORC in the early 2000s (110). Evanescent and cross-couplers and Bragg gratings were first demonstrated at 1550 nm (111; 112), and the waveguides were integrated with microfluidic channels as chemical sensors (113). From the start the new technology appeared very promising, with most reported losses below 0.5 dB/cm (114) and some below 0.2 dB/cm (115); there are reasons to believe that this can be much improved in the near future given the renewed interest in the technique (104).

Until quite recently, these UV-written waveguides were designed for operation at 1550 nm, primarily for telecom applications (splitters, Bragg gratings filters, multiplexers, add-drop etc.) With the explosion of interest for lab-on-a-chip devices and biochemical sensing applications in recent years, the ORC started to work on waveguides operating in the visible/near-infrared regime because the sensitivity of such devices at 1550 nm is compromised by the strong absorption of OH bonds in water (in excess of 30 dB/cm above 1400 nm (116)). It was in relation with the work presented in Chapter 5 that the first experiments with 780 nm UV-written waveguides were reported. Moving from 1550 to 780 nm requires significant changes to the apparatus to focus the UV beam more tightly (since a smaller core is required), which also reduces significantly the tolerance to alignment errors. The propagation losses at 780 nm were virtually unknown.

The writing process is illustrated in Fig. 2.3. First, the chip itself is made by standard deposition methods (see previous section). A first layer of thermal oxide (16 μm) is grown on the silicon substrate. Then a germanium and boron-doped core silica layer (5.6 μm), followed by an overclad silica layer (17 μm), are deposited by FHD. The germanium doping of the core layer provides photosensitivity, while boron is used to match the refractive index of the core layer to that of the adjacent layers. For the writing itself, a UV laser beam (frequency-doubled ion-argon, $\lambda = 244$ nm) is tightly focused to a single spot on the active layer, and any structure can be written by translating the substrate. The actual apparatus uses a double-beam set-up that lets us write Bragg gratings as well, as detailed in Chapter 5.

An interesting and novel device made possible by direct laser writing is the X-coupler, fabricated simply by crossing two waveguides. These were previously difficult to make because the standard deposition/etching technique does not allow one to make sharp wedges and corners, especially with a shallow angle of a few degrees. Devices that allow energy exchange between two waveguides were limited to directional couplers, where two guides are brought in close proximity so that their mode fields overlap and the evanescent field of one waveguide then resonantly excites a mode in the other. Although X-couplers operate on a similar principle, as detailed in section 2.4.2, they can be much more compact for a given coupling ratio. They can be tuned by varying parameters such as the crossing angle or the relative fluence of the two crossing guides. Maximum and minimum coupling ratios of 95% and 2% have been experimentally achieved with angles smaller than 5°, with relatively little wavelength or polarisation dependence (111).

Another very promising device is the so-called flat fibre (104). Starting with a preform similar to classical optical fibres but with a modified drawing process, they have a flat photosentitive core on which waveguides can be written. The essential advantage of these devices is that they can potentially achieve very low losses, similar to optical fibres (0.2 dB/km at 1550 nm and 3 dB/km at 780 nm). They are still, however, in the very early stages of development.

2.2.3 Femtosecond laser writing: Jena chip

Femto-second writing in bulk silica (117) was invented a few years after the direct UV-write method in silica-on-silicon chips. The two have been progressing in parallel since then. A femtosecond pulsed infrared laser (usually Ti:Sapph, 800 nm, ~ 50 fm) is focused inside a piece of bulk glass. Absorption at this wavelength is normally negligible, but the very high peak power allows non-linear, multiphoton absorption leading to melting under the focal spot. The glass is denser after re-solidifying, leading to an increased refractive index with a graded profile (118). Waveguides with arbitrary 3D geometries can be written by translating the sample under the focus (105). The process does not

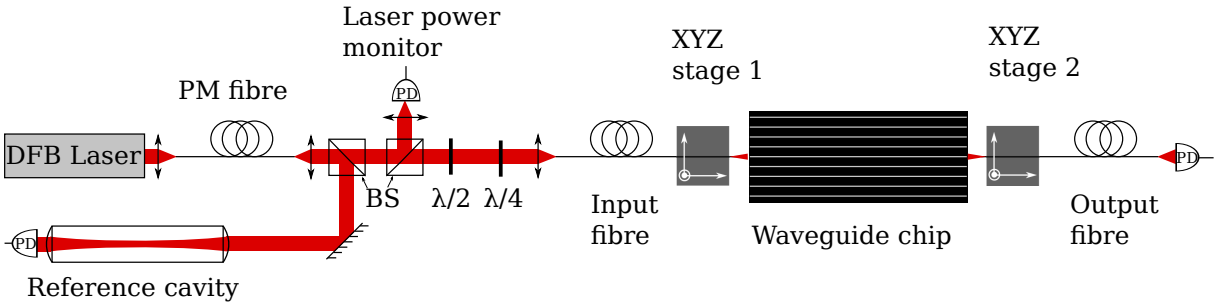

Figure 2.4: Experimental set-up for coupling in and out of the waveguide.

depend on the intrinsic photosensitivity of the substrate, and it is not limited to fused silica. It is even possible to write waveguides inside gain materials to create amplifiers and lasers (119; 120) or in electroactive materials such as lithium niobate (121).

All the standard devices have been demonstrated using this technique (splitters, Mach-Zehnder interferometers, directional couplers, Bragg gratings..., see references above). Low-loss waveguides with high uniformity can be achieved by writing the device in multiple passes so that the defects of single passes cancel out; typical propagation losses are around 0.2 dB/cm at 1550 nm, comparable to the the direct UV-writing technique. The index contrast can be as high as 0.01. The use of adaptive optics can further improve the resulting waveguide by getting rid of optical aberrations by compensating the spherical aberration introduced by the air-glass interface. Femto-second written waveguides have already been used in quantum optics experiments (122).

2.3 Waveguide characterisation

The most relevant properties of the waveguides are, for our purpose, the mode profile and the propagation loss. The latter will be addressed mainly in the next chapter, so we focus here on the waveguide's mode profile, which we would like to match that of a standard 780 nm optical fibre (Nufern 780-HP). We assume that the waveguide and fibre modes can be well approximated by Gaussian modes.

2.3.1 Methods

The measurement set-up is sketched in Figure 2.4. A 780 nm laser diode is coupled to a PM fibre, and then into a standard 780 nm single mode fibre mounted on a three-axis flexure stage to be aligned to the desired waveguide. On the other side of the chip, a second fibre is aligned with the waveguide and fed to a photodiode. For transmission measurements, an index-matching oil is added between the chip end facet and the fibres to eliminate Fresnel losses. The reference cavity allows tracking of the laser frequency and is not used here.

The field transmission $\eta(\Delta x, \Delta z)$ between fibre and waveguide as a function of misalignment is determined by the convolution of their respective modes, as detailed in Appendix A, Eq. (A.6). This is maximum when they are aligned ($\Delta x = \Delta z = 0$), resulting in an intensity coupling efficiency η^2 (Eq. (A.7)). The total transmission is then $T = \eta^4 \exp(-L_{\mathrm{dB}} l/10)$, with l the waveguide length in cm and L_{dB} the propagation loss in dB/cm.

According to Eq. (A.6), the mode size of the waveguide $w_{x,z}$ (defined as the $1/e^2$ half-width) can easily be determined if the fibre waist w_f is known. One can record transmitted intensity while scanning the fibre in front of the waveguide, and then deconvolve the measured width $W_{x,z}$ to recover $w_{x,z}$ as $w_{x,z} = \sqrt{W_{x,z}^2 - w_f^2}$. The fibre mode size can be measured similarly by scanning two identical fibres across each other; then $w_f = W_f/\sqrt{2}$.

This method is in principle direct, simple and convenient, but in practice the interpretation of the results is fraught with problems. The first is the necessity to bring the scanning fibre close enough to the waveguide output that the beam expansion can be neglected. With a typical Rayleigh range of order $25\,\mu$m, this requirement translates in distances measured in microns, when the positioning accuracy under a microscope is no better than $5 - 10\,\mu$m.

The second difficulty is the calibration of the fibre displacement. In the early stages of this experiment, we only had at our disposal a three-axis flexure stage whose piezo actuators did not include position sensors. Moreover the piezo displacement was highly non-linear. To calibrate it, we imaged the output of a fibre mounted on the stage on a CCD and measured its centroid with sub-pixel resolution. The same voltage driving sequence was used for the calibration and the actual measurement.

The second method is a far-field measurement. The mode at the waveguide output is left to expand freely for a centimetre or so before being imaged on camera at several distances within a coupled of centimetres. We can then measure the beam divergence $\theta = \lambda/\pi w_0$. This straightforward method has its own drawbacks: it requires a cleanly polished output facet, which was not the case on all of our chips. Moreover the camera had a protective glass window which introduced interference patterns, adding uncertainty to the mode size measurement. This problem can be partially alleviated by operating the laser below threshold: the large bandwidth of amplified spontaneous emission washes out the interferences.

A third method is to image the mode with a high-power immersion microscope objective whose magnification can be accurately calibrated. It may be asked whether that calibration, carried out on a geometric image, remains valid for a Gaussian beam. This is not a trivial question, and we show that it is the case in Appendix A.3.

In the following experiments we used all of these methods in various measurement carried out over several years.

2.3.2 ORC chips

We wrote a UV-written chip ("ORC1") containing 16 waveguides written with fluence F varying between 4 and 19 J/cm^2. The mode profiles were measured with the fibre scan method and are diplayed in Figures 2.5 and 2.6(b). The vertical mode size w_z is constant whereas the transverse one w_x has a clear minimum around 14 J/cm^2 (see geometry in Figure 2.6(a)). We interpret this behaviour as follows. In the vertical direction, the index contrast is dominated by the mismatch between the core and lower/upper cladding layers, so the mode size is relatively insensitive to the fluence and has an almost square profile. Transversely, the index profile depends on the laser intensity profile and on the index change vs fluence characteristic, which is linear at low fluence but saturates at high power (see Fig. 5.5, page 101). At low fluence, i.e in the linear regime, the index profile is Gaussian, following the UV laser intensity profile, and the weak index contrast results in a large mode which gets smaller as the fluence and the index contrast increases. The waveguide dimension however remains constant before the onset of saturation. At higher fluence, the index change saturates, producing an index profile that is flattened out – effectively resulting in a larger waveguide so that the mode size increases again.

On a second chip ("ORC2"), we wrote 5 sets of 10 waveguides, each set having a fluence of 12, 13, ..., 16 J/cm^2. We measured the mode profiles again and confirmed that the mode is reasonably circular for a fluence around 14 J/cm^2, with $w_x \times w_z = 4.3(5) \times 3.0(5)\,\mu$m^2, from which we estimate a maximum mode-matching efficiency $\eta_{\mathrm{w}}^2 = 90 \pm 5\%$, the fibre mode size being $2.6 \pm 0.3\,\mu$m. We also measured the transmission of all these waveguides, which is close to 50% with a very large dispersion ($\pm 10\%$) due to differences in coupling efficiencies from shot to shot. We can compare this to the transmission expected from the mode profile by using Eq. (A.7) and taking into account a propagation loss of about 0.8 dB/cm (this will be determined in the next chapter). This yields a theoretical transmission of $60 \pm 10\%$, quite compatible with the measured transmission.

We did not attempt to measure directly the index profile of the waveguides, and therefore did not carry out precise simulations using MPB since we only care about the mode itself. The effective index of the waveguides was measured at the ORC using Bragg gratings: $n_{\mathrm{eff}} = 1.4590(5)$.

2.3.3 Jena chip

This femto-second written sample was the only one of its kind we were able to secure. As depicted in Figure. 2.7, it contains 20 sets of 8 identical waveguides. Each set was written with a different combination of laser power and writing speed. Higher power and slower speed correspond to higher index contrast, although these two parameters are not strictly equivalent and they cannot be characterised by a single parameter like the fluence we used in relation to UV-written waveguides. Within a set, the waveguides are written

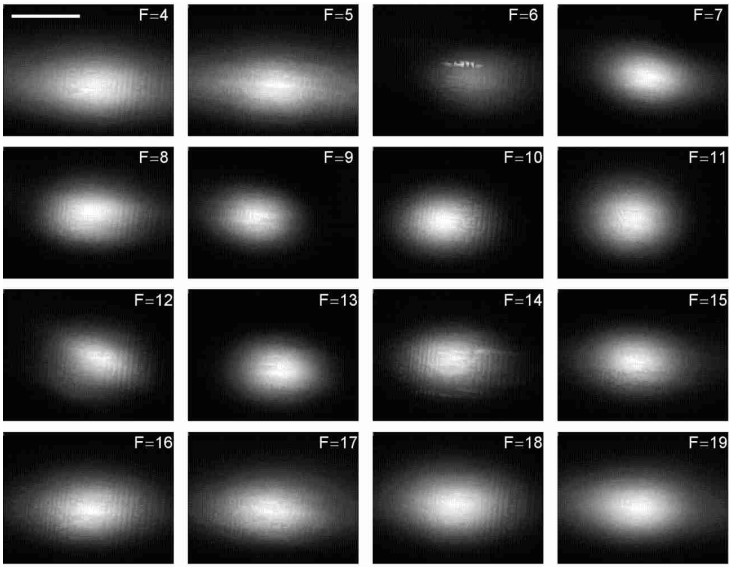

Figure 2.5: Mode profiles on chip ORC1. The writing fluence (in J/cm^2) is indicated on each image. The scale is indicated by the 5 μm long white line in the first image.

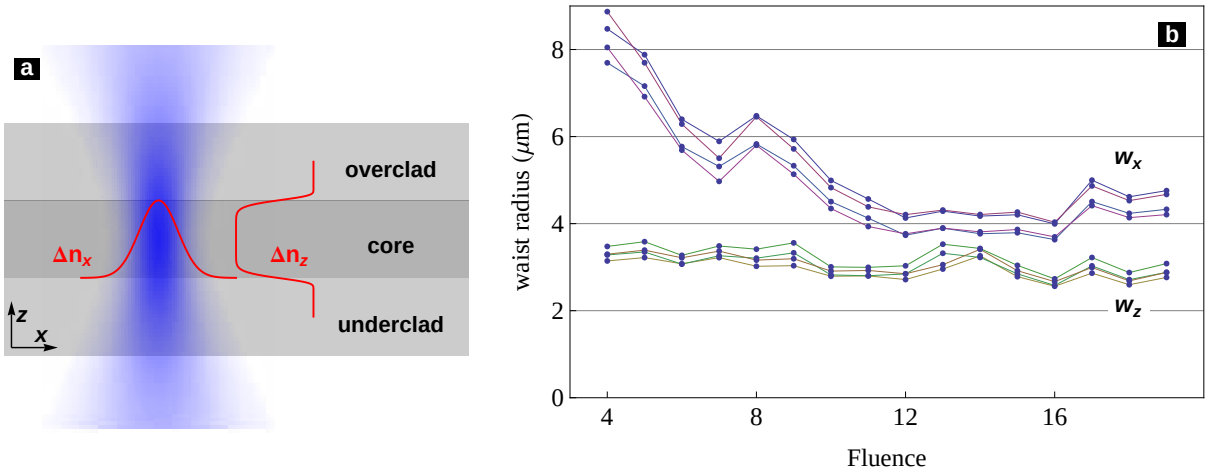

Figure 2.6: Mode profiles on chip ORC2. **(a)** Writing geometry, showing the photosensitive core layer, the UV laser beam and the typical index profiles. **(b)** Mode size as a function of writing fluence (in J/cm^2). The waist in the vertical direction w_z is essentially constant, being determined by the thickness of the Ge-dopped layer, which is not perfectly index-matched to the lower and upper cladding layers. The horizontal waist w_x is very large at low fluence and decreases as the waveguide index contrast increases and the light is more tightly confined. w_x increases again at large fluence, where the waveguide becomes larger due to fluence saturation. The four curves were obtained from the forward and backward raster scan and for two different calibrations of the piezo displacement.

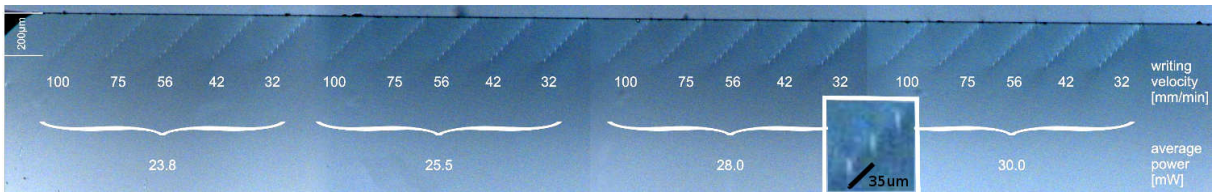

Figure 2.7: Photograph of the output facet of femtosecond-written Jena chip. It consists of 4×5 sets of 8 identical waveguides, each set written at a different fluence by changing both the sample translation speed and the laser power. Unlike UV-writing, the two parameters are not strictly equivalent. The inset shows a magnified detail, highlighting the waveguide's asymmetry. Image courtesy of Jena University.

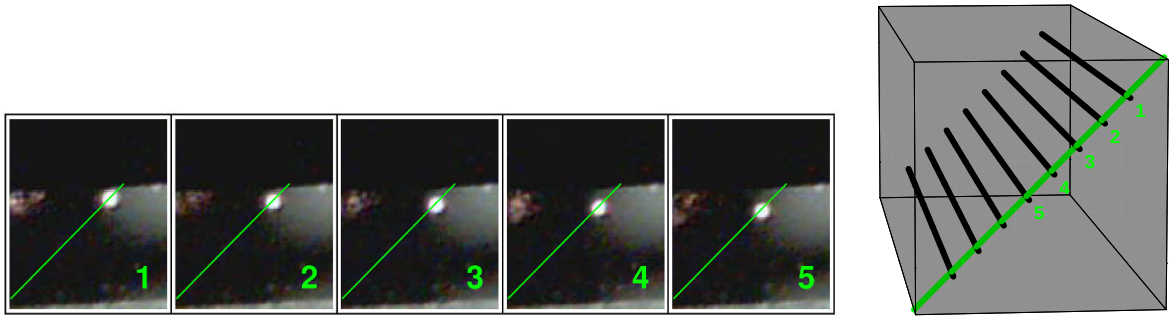

Figure 2.8: Microscope images of the output of one set of waveguides on the Jena chip, as the in-coupling fibre is moved between the first 5 waveguides. The sketch on the right clarifies the geometry. The green line is a guide for the eye indicating the position of the waveguides.

at increasing depths in steps of 25 μm. Figure. 2.8 shows the waveguide output as the input fibre is moved between different waveguides.

We measured the mode profile and transmission as before. The results are summarised in Figure. 2.9. The modes have a pronounced asymmetry which is already visible in the microscope image in Figure. 2.7 (see inset). As expected, higher laser power and slower writing speed results in smaller mode because the induced index contrast is larger. Smaller modes then result in higher transmission by virtue of their better mode-matching to the input and output coupling fibres. There is a reasonably good agreement between the measured transmission and that expected from the mode size, even without including the propagation loss. The latter was measured in the next chapter and found to be 0.3 dB/cm or about 10%, which is indistinguishable from zero within the measurement error.

2.4 Couplers

We have already introduced these devices. In this section we first review the coupled-mode theory for directional couplers, focussing on the derivation of the relevant parameters (especially the coupling constant and coupling ratio); the theory is dealt with in more detail in many textbooks (see for example (98)). We then describe ORC's cross- and evanescent couplers before presenting experimental measurements of the coupling ratios.

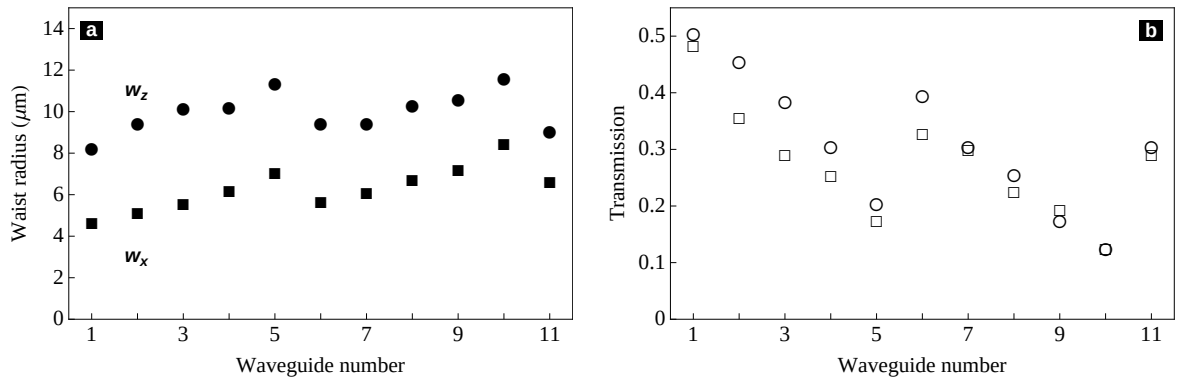

Figure 2.9: Mode profiles on the Jena chip. **(a)** Measured spot size of the first 11 Jena waveguides (note: this is the raw spot size, before deconvolution of the fibre mode). Notice the pronounced ellipticity. **(b)** Circles: measured transmission of the same waveguides. In the first 5, the laser power was 30 mW and the writing speed increases from 32 to 100 mm/min, resulting in weaker index contrast and increasingly larger modes, and so a decrease in transmission. For the next five, the laser power was reduced to unit 28 mW, resulting in slightly larger modes. The transmission calculated from the mode sizes via Eq. A.7 is plotted as squares. They are consistent with the measurement given error bars of ±5%.

2.4.1 Theory

We start by considering two identical waveguides, separated by a distance d, and we denote the unperturbed mode field functions by $E_{1,2}$. They are normalised according to $\iint_{-\infty}^{\infty} |E|^2 d\mathbf{r}^2 = 1$ and propagate with the same k-vector $\beta_0 = kn_{\text{eff}}$. $k = 2\pi/\lambda$ is the vacuum wave vector and n_{eff} is the waveguide's effective index. The second waveguide is considered as a polarisation perturbation from the point of view of the first one, which leads to a cross-coupling coefficient κ_l and a self-coupling coefficient M_l (98):

$$\left\{ \begin{array}{rcl} \kappa_l & = & \frac{k^2}{2n_{\text{eff}}} \iint_{-\infty}^{\infty} \Delta\epsilon E_1 E_2 d\mathbf{r}^2 \\ M_l & = & \frac{k^2}{2n_{\text{eff}}} \iint_{-\infty}^{\infty} \Delta\epsilon |E_1|^2 d\mathbf{r}^2 \end{array} \right. \tag{2.5}$$

and to the coupled equations

$$\left\{ \begin{array}{rcl} dE_1/dz & = & -i(\beta_0 + M_l)E_1 - i\kappa_l E_2 \\ dE_2/dz & = & -i\kappa_l E_1 - i(\beta_0 + M_l)E_2 \end{array} \right. . \tag{2.6}$$

The self-coupling coefficient M_l expresses the modification of the modal propagation constant due to the proximity of the other waveguide, and is often experimentally irrelevant. κ_l is the cross-coupling coefficient, in units of m^{-1}, and characterises the number of coupling cycles per unit length.

Since $M_l \ll \beta_0$ we neglect it, and Eq. 2.6 is easily solved under the initial condition $E_2(0) = 0$ (power incoming from the first waveguide only). We find that the intensity

along the coupler evolves as:

$$\begin{cases} I_1(z) &= I_1(0)\cos^2\kappa_l z \\ I_2(z) &= I_1(0)\sin^2\kappa_l z \end{cases} \tag{2.7}$$

So the light is coupled back and forth between the two waveguides; the length required for total coupling is

$$I_2(L_c) = 1 \Rightarrow L_c = \pi/2\kappa_l \tag{2.8}$$

We have assumed throughout that the two waveguides are identical, so they are phase matched ($\Delta\beta = \beta_2 - \beta_1 = 0$). If this were not the case the maximum coupling ratio would be smaller than unity.

The system of equations (2.6) can be written in matrix form and it is then straightforward to generalise to linear systems of coupled waveguides, as illustrated for example in Fig. 2.10. Under the assumption of nearest-neighbour coupling, we write the fields in waveguides $1, ..., N$ as a column vector $\mathbf{E}(z) = \{c_n(z)\exp(-i\beta_0 z)\}_N$, and it follows that:

$$\frac{d\mathbf{E}}{dz} = -i\mathbf{C}\mathbf{E}, \quad \mathbf{C} = \begin{pmatrix} \beta_0 + M_l & \kappa_l & 0 & \cdots & \cdots \\ \kappa_l & \beta_0 + M_l & \kappa_l & 0 & \cdots \\ 0 & \kappa_l & \ddots & \ddots & 0 \\ \vdots & 0 & \ddots & \ddots & \kappa_l \\ \vdots & \vdots & 0 & \kappa_l & \beta_0 + M_l \end{pmatrix}. \tag{2.9}$$

The eigenvectors of this matrix define propagating normal modes, that is modes of the coupled array that are invariant under propagation. $\mathbf{C}$ is easily diagonalised, and each of its N eigenmodes (denoted by the subscript $s = 1 \ldots N$) has a unique propagation constant β_s given by

$$\beta_s = \beta_0 + M_l + 2\kappa_l \cos\frac{s\pi}{N+1} \tag{2.10}$$

while the normal modes will consist of a superposition of the unperturbed modal field with weights given by the eigenvectors

$$a_{s,n} = \sin\left[\frac{sn\pi}{N+1}\right]. \tag{2.11}$$

The eigenmodes in the case $N = 5$ are plotted in Figure 2.10.

Calculating the coupling coefficients

The coupling coefficients can be computed directly from Eq. (2.5) and from the knowledge of the waveguides index and mode profile. As an alternative to evaluating these overlap integrals, we use MPB to compute directly, the propagation constants β_s of the two normal

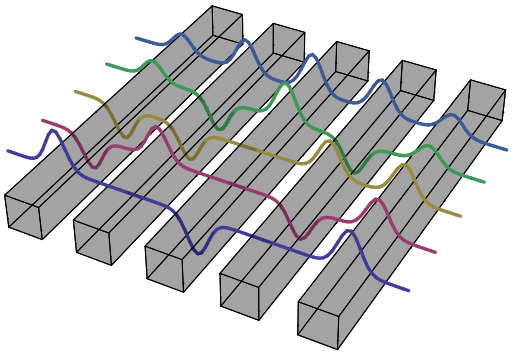

Figure 2.10: An array of five coupled waveguides and the associated normal modes.

modes in the $N = 2$ case. Then from Eq. (2.10) we see that

$$\kappa_l = \frac{\beta_1 - \beta_2}{2} \tag{2.12}$$

$$M_l = \frac{\beta_1 + \beta_2}{2} - \beta_0 \tag{2.13}$$

Since the overhead of calculating β_1 and β_2 is very small once we have set up the calculation for a single waveguide, this method is very straightforward.

Analytical solution From the solutions to the cylindrical waveguide problem presented earlier in section 2.1, we can derive an analytical expression (originally due to Snyder (123)) for the cross-coupling coefficient κ_l, in the weak-guiding approximation:

$$\kappa_l = \frac{1}{a^2 n_{\text{co}} k} \frac{u^2}{v^2} \frac{K_0(wd/a)}{K_1^2(w)} \tag{2.14}$$

where as before a is the waveguide radius and d the distance between the centres.

Fig. 2.11 shows the evolution of κ_l as a function of the separation between two square waveguides, using the three methods just discussed. We can see that the agreement is very good, even though the analytical solution is actually for cylindrical fibres; this is not so surprising since the modes of cylindrical as well as square waveguides are well approximated by a Gaussian.

For example, for a mode of Gaussian waist $w = 3\,\mu$m, $\Delta n = 4 \times 10^{-3}$, and $d = 8\,\mu$m (parameters typical for our waveguides) we get $\kappa_l = 0.92\,\text{mm}^{-1}$, and a 0.9 mm-long evanescent coupler results in a coupling ratio of 54%.

2.4.2 ORC's X-couplers

So far we have only talked about traditional evanescent couplers. While they are widespread devices, they suffer from a typically low coupling strength κ_l, resulting in mm-scale cou-

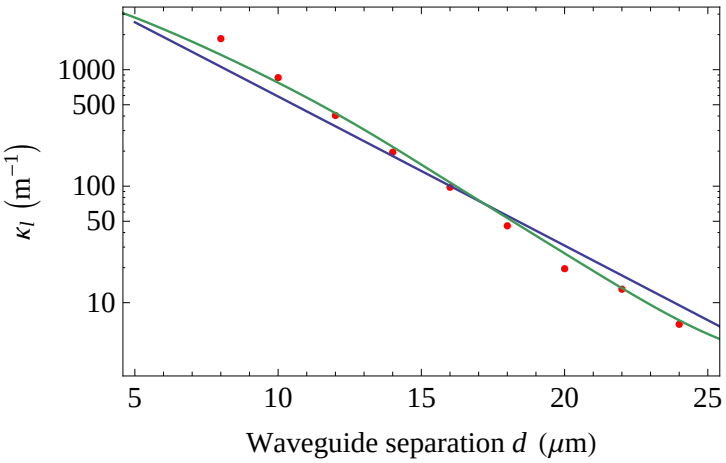

Figure 2.11: Coupling constant κ_l as a function of waveguide separation d, according to the three methods outlined in the text. Green curve: integral overlap method for a 5 μm square waveguide, calculated with MPB. Red dots: direct MPB simulation (Eq. (2.12)), with the same parameters. Blue curve: Snyder's analytical formulae Eq. (2.14) for a cylindrical fibre of diameter 5 μm. The index contrast was $\Delta n = 0.004$.

plers, to which the lengths of the bends required to bring the waveguides together must be added. An alternative is to cross the two waveguides at a shallow angle of a few degrees such that the crossing region supports at least the first two symmetric and antisymmetric modes. The input TEM_{00} mode can be decomposed in terms of higher order modes (normally only two) which propagate with a different k-vector, so that the relative amplitudes in the two single-mode outputs depends on the length of the device and on these k-vectors, which are also functions of the crossing angle. The operation is then identical to that of evanescent couplers. The main advantage over evanescent couplers is the lower wavelength and polarisation dependence, as well as a higher reproducibility which result from the much larger mode coupling and shorter device length.

Geometry

The X-couplers developed at ORC are an original compact design (111), as shown in Figure 2.12(a). It involve four interrupted cosine bends (in the grey-shaded area), joined by two straight lines intersecting at a full angle α. We define a cartesian coordinate system $\{x, y\}$ on the chip, with waveguides running along x. The basic cosine S-bend of length l and height $d/2$ is parametrised as

$$y_b(x) = \frac{d}{2}\left[1 - \cos\left(\pi\frac{x}{l}\right)\right]. \tag{2.15}$$

This is the equation for the bottom left bend in Figure 2.12(a); the three others have very similar expressions. The full angle α at any point along the bend is given by $\tan\alpha/2 =$

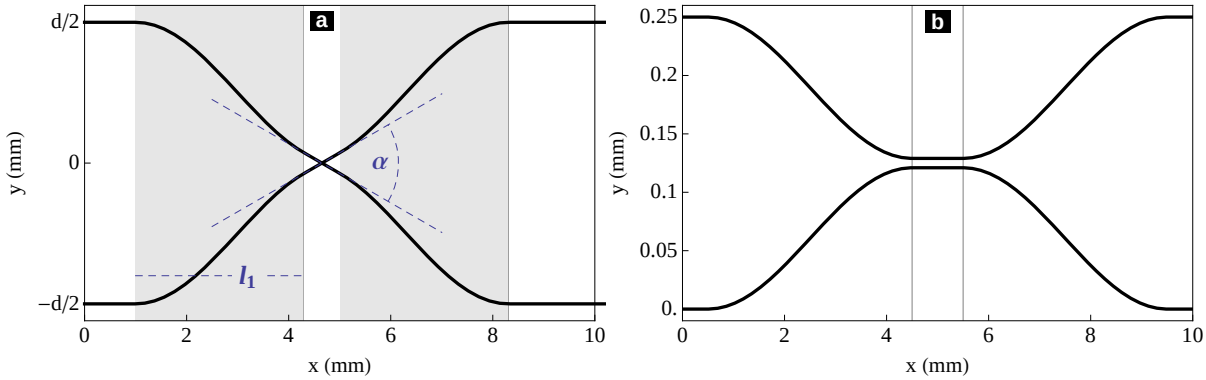

Figure 2.12: Typical cross- and evanescent coupler. The vertical lines indicate the transition between cosine bends and straight sections. The angle of the X-coupler is 3°.

dy_b/dx, hence

$$\alpha(x) = 2\arctan\left[\frac{d\pi}{2l}\sin\left(\pi\frac{x}{l}\right)\right] \tag{2.16}$$

For typical bend parameters $l = 4\,\text{mm}$ and $d = 0.125\,\text{mm}$, the maximum angle at the centre of the bend is 5.6°. To get the desired angle, one has to interrupt the S-bend after a length l_1 which is found by solving the previous equation for l:

$$l_1(\alpha) = \frac{l}{\pi}\left(\pi - \arcsin\left[\frac{2l\tan(\alpha/2)}{d\pi}\right]\right) \tag{2.17}$$

The truncated S-bend has a height $d_1 = y_b(l_1(\alpha))$. It needs to be completed with a straight line to maintain the total height d; this line must run for a distance $l_2 = (d-d_1)/\sin(\alpha/2)$.

We now can use these equations to draw a complete cross-coupler, as depicted in Figure 2.12. The larger the angle, the shorter the total length of the coupler will be, although even the steepest one will still require about 7mm. Given the need to integrate heaters on both sides of the couplers, the minimum chip length will be about 10mm.

Operation and optimisation

There are two main sources of geometric losses in waveguides: bends, which have radiation losses due to the curvature, and junctions between a bend and a straight section, which have loss due to the curvature mismatch. The latter is essentially eliminated by the use of cosine bends instead of the traditional circular arcs. Many bends were fabricated at the ORC and their losses measured; it was found that for a bend height of 125 μm, S-bends could be made as short as 4 mm before the loss became measurable. The upper bound on single bend loss is then 0.1 dB (111).

There is of course no reason why the index contrast resulting from writing twice, with the same fluence, over the coupling region should result in waveguide with the appropriate multimodal behaviour. Moreover we need to compensate for the "proximity effect", whereby the photosensitivity is altered in the vicinity of a waveguide that has

already been written. The X-couplers were optimised at the ORC for minimum losses and symmetric behaviour (i.e. the coupling ratio should be independent of which input waveguide is used) by varying the fluence (the writing speed) of both passes. It was found that, for a base fluence (in the horizontal sections and S-bends) of $17\,\mathrm{kJ/cm^2}$, the optimum fluence for the first straight junction was $12\,\mathrm{kJ/cm^2}$, and $18\,\mathrm{kJ/cm^2}$ for the second one (111). All X-couplers in this work were written with these known parameters, although $17\,\mathrm{kJ/cm^2}$ does not produce the most circular modes. A new round of optimisation would be required for X-couplers based on $14\,\mathrm{kJ/cm^2}$ waveguides.

2.4.3 Evanescent couplers

We also produced at the ORC a series of evanescent couplers, based on the same optimised S-bends. The geometry is shown in Figure 2.12. The main motivation for using evanescent couplers is the expectation of lower loss compared to X-couplers. The latter indeed suffer from scattering caused by the relatively sharp transition between the single waveguides and the coupling region, which is a bit short of adiabaticity. The coupling ratio can in principle be tuned by varying the distance between the waveguides and/or the length. However earlier (unpublished) work at the ORC on evanescent couplers was inconclusive: the coupling ratios exhibited large variations from the expected trends in addition to a lack of reproducibility between nominally identical couplers. But since we intend to tune the coupling with phase shifters, this isn't a major problem.

2.4.4 Our devices: coupling ratios measurements

In this section we describe measurement of the coupling ratios, on different chips and with a range of methods. The coupling ratio is defined as $C_\times = I_\times/(I_\times + I_=)$, where $I_=$ is the power detected at the output of the input waveguide, and $I_\times$ the power detected at the output of the second waveguide (the notation is made clear by the inset sketches in Figure 2.13). Although straightforward in principle, coupling ratio measurements are fraught with difficulties in principle.

Methods and issues

Method 1: The first method consist in outcoupling light with a fibre and index-matching oil. A single output fibre is used, which is moved between the two output as required, and fed to a photodiode. This method suffers from large variations in waveguide-fibre coupling efficiency from one waveguide to the next, and even between measurement on a single waveguide. It was used on samples ORC4 and ORC6. The coupling problem can be alleviated to some extent with the following consideration. Let us define the two output voltages as $\{V_1, V_{12}\}$ when the input fibre is coupled to the first waveguide, and

$\{V_{21}, c_{22}\}$ when it is coupled to the second waveguide. η_1, η_2 are the (unknown) coupling efficiencies between the two output waveguides and the out-coupling fibre. Assuming that the coupler is symmetrical, the ratio C' of the two outputs must be independent of the input:

$$C' = \frac{I_\times}{I_=} = \frac{V_{11}/\eta_1}{V_{12}/\eta_2} = \frac{V_{22}/\eta_2}{V_{21}/\eta_1} \Longrightarrow C' = \sqrt{\frac{V_{11}V_{22}}{V_{12}V_{21}}} \tag{2.18}$$

The cross-coupling ratio $C_\times = I_\times/(I_\times + I_=)$ is $C_\times = 1/(1 + C')$.

Method 2: A second method, used on sample ORC6 and ORC7, is to image the outputs on a CCD using a microscope objective. With a low-power objective (10x), both outputs can be imaged at the same time. It requires that the chip output facet be very well polished: in addition to scattering, surface roughness imparts onto the beam waist a phase profile which results, on the CCD, in complicated interference patterns overlaid over the Gaussian mode. This makes estimating the beam radius difficult and inaccurate. The intensity I is obtained by integrating over a square window centred on the mode, whose size is chosen to avoid clipping the mode while minimising background. Because the dynamic range of the camera is only 8 bits, we also change the exposure time t_{exp} independently for each output so that the peak intensity in the integrating window is safely below saturation, and we record the ratio I/t_{exp} as the optical intensity.

Method 3: The last method, which we used on the final samples (ORC8), is to use an immersion high-power objective (100x). The index-matching oil almost cancels any scattering and additional random phases, even on these unpolished samples. On the other hand only one mode can be imaged at a time, which makes the measurement somewhat slower. The same consideration about changing the exposure time apply. This was found to be the most accurate method.

Other problems: The coupling ratio measurement was made more difficult by a few other technical challenges. Most importantly, the second mode of these waveguides operating at 780 nm is quite close to the cut-off frequency, and in fact quite a few of them seem to be able to support it as their mode sometimes exhibits side lobes. This makes the coupling ratio dependent on the input fibre alignment, which changes the relative amount of input power between the different modes. Unfortunately, optimising the input fibre position to maximise the output power does not select the first-order mode, but a combination of both. This difficulty is the main cause of uncertainty.

Results

We wrote several generations of couplers on a number of chips. We only present here results from the last generation (ORC8). Four nominally identical chips were written,

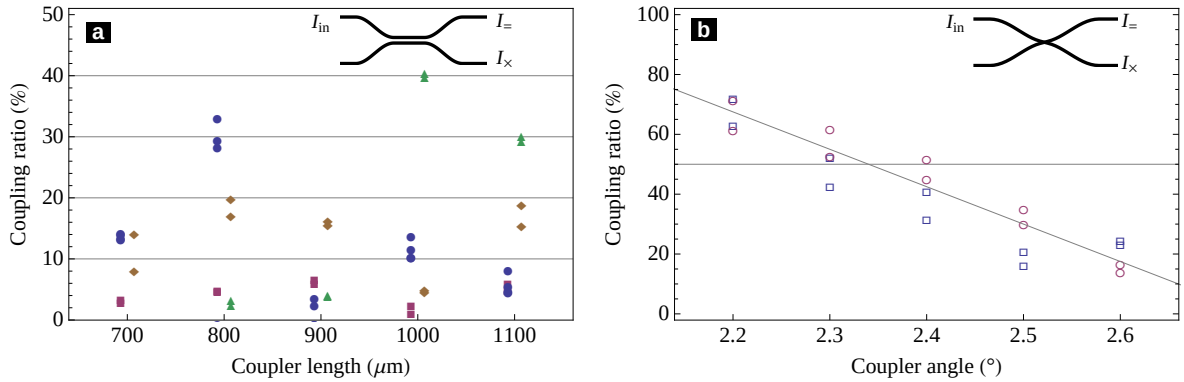

Figure 2.13: Measured coupling ratios of 4 sets of nominally identical evanescent couplers (a) and 2 sets of (also identical) cross-couplers (b). Markers colour/shape indicate different sets. The coupling ratios were measured from both inputs and both are plotted, to give an idea of the couplers' symmetrical behaviour. The trend line in (b) is only a guide for the eye.

two on a legacy wafer (W7) and two on a new wafer (W10), which was expected to have lower losses. Each chip contained one set of five cross couplers and another set of five evanescent couplers, although a programming error made two sets of cross-couplers unusable (so we had 20 evanescent and 10 cross-couplers in total). The coupler lengths or angles were chosen to obtain a range of coupling ratios around 50% (0.7 to 1.1 mm and 2.2 to 2.6°). The measurements were performed according to Method 3 (100x immersion objective).

Results for the evanescent couplers are plotted in Figure 2.13(a). There are eight data points for each coupler lengths, since for each of the four identical couplers on the four chips we measured $C_\times$ for the two inputs. Two facts are immediately apparent. First there is quite a large dispersion in the measured coupling ratios for a given coupler length, spanning almost 40% in three out of five cases. This is in line with the poor reproducibility of these devices previously observed at the ORC. On the other hand the coupling ratio is relatively independent (within $\sim 5\%$) of the input waveguide, indicating good symmetry. Secondly, we do not observe any clear trend, whereas we expected the coupling ratios to increase with the coupler length from 40 to $\sim 60\%$. It is possible that the trend is simply buried in the large dispersion of the data.

We now turn to the cross-couplers in Figure 2.13(b). For each angle, the four data points correspond to two identical couplers on two different chips (red circles/blue squares), measured from the two inputs. We notice a modest ($\sim 10\%$) but consistent asymmetry between the two inputs, together with a similar variation between identical couplers. The coupling ratio decreases clearly from $\sim 65\%$ for a crossing angle of 2.2° to $\sim 20\%$ at 2.6°, in line with previous experiments as well as theoretically calculations performed at ORC (111).

2.5 Thermo-optic phase shifters

2.5.1 Introduction: electro- vs thermo-optics phase shifters

Phase shifters on integrated optics devices rely either on the thermo-optic or electro-optic effects. The latter has become mainstream technology in telecom devices thanks to its high bandwidth (40 GHz and beyond), whereas the former is much slower (kHz) but easier to manufacture.

The electro-optic or Pockels effect is a second-order, $\chi^{(2)}$ non-linear effect present only in crystalline materials without inversion symmetry. As such, it does not exist in fused silica unless some special treatment known as thermal poling is applied (124); even then the achievable phase shift is too small for practical applications. A common material for electro-active waveguides is lithium niobate ($LiNbO_3$). Gigahertz bandwidth are now common. However such waveguides are complex and expensive to manufacture, and since we do not require a large bandwidth we settled for the thermo-optic alternative.

The thermo-optic effect relies on the fact that the refractive index of silica (and of other optical materials) is a function of temperature. The thermo-optic coefficient of fused silica is in the range $dn/dT = (0.9 - 1.2)10^{-5}\,\mathrm{K}^{-1}$ depending on the glass composition (125), and the induced phase shift is given by

$$\Delta\phi = k_0 l \Delta n = k_0 l \left(\frac{dn}{dt}\right) \Delta T. \tag{2.19}$$

A phase shift of π over a length $l = 1\,\mathrm{mm}$ then requires a temperature change of only 15 K. The heat can be provided by small resistors glued to the chip, or by directly depositing a strip of resistive material, both of which are comparatively simple. Thus thermo-optic phase shifter seem ideally suited for our purpose. They have become increasingly popular in relation to silicon photonics and are also being used in on-chip all-optical QIP (Bristol group) (126).

2.5.2 Fabrication of thermo-optic phase shifters

We fabricated our own phase shifters in the CCM coating plant. A 400 nm layer of 20/80 nickel-chromium alloy (NiCr, or Nichrome), a common high-resistivity material, was sputtered on glass slides for testing, and eventually on the waveguide chips. Gold wires and mm-wide pads were subsequently deposited to connect the wire [2] , as in the schematics of Figure 2.14, which shows the final NiCr mask design for 3 couplers (4 heaters

[2] The deposition recipe is as follows:
NiCr: DC plasma, strike at 10 W/15 SCCM, coating at 110 W/9 SCCM. Deposition rate 0.3-0.4 Å/s.
Gold: RF plasma, strike at 5 W/10 SCCM, coating at 130 W/6 SCCM. Deposition rate 0.5-0.6 Å/s.
The base pressure in the vacuum chamber is below 5×10^{-7} Torr. The initial thickness target was 200 nm, and resulted in 400 nm-thick layers, probably due to the monitor crystal calibration. We decided to stick with 400 nm.

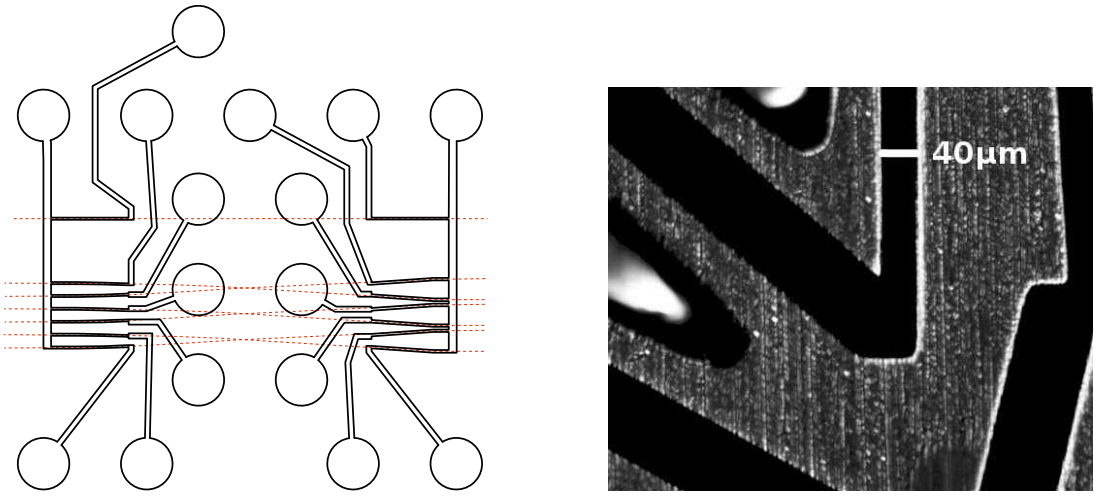

Figure 2.14: Left: Final NiCr mask design with two heaters positioned on top of a straight waveguide, and three sets of four heaters on all four arms of three different couplers (waveguides are indicated by the dashed red line). Right: Microscope image of the mask. The width of the slit is 40 μm.

Figure 2.15: Microscope images of sputtered wires. Left, NiCr layer only on a glass slide. Right, NiCr and gold on a waveguide chips. The waveguides can be seen emerging from underneath the heaters. The rugged area to the left is some residue from the mirror coating that has spilled onto the chip top surface.

each) and a single straight waveguide (2 heaters). The gold mask (not shown here) would look very similar, except that the narrowest sections are removed so as not cover the NiCr with gold.

We initially attempted to fabricate the masks on the CCM workshop wire eroder, but the smallest slit dimension was limited to about 180 μm. We then subcontracted the mask fabrication to Tecan Ltd, whose photo-chemical etching process offers resolution down to 10 μm. A micrograph of the final stainless steel mask with 40 μm slits is also shown in Figure 2.14.

Microscope images of a NiCr layer deposited on a glass slide, and of NiCr and gold layers on a waveguide chip, are shown in Figure 2.15. Figure 2.16(a) shows typical AFM profiles for two wires made with 40 and 50 μm masks. The actual width of the wire is identical to that of the masks, and the thickness is consistently about 400 nm.

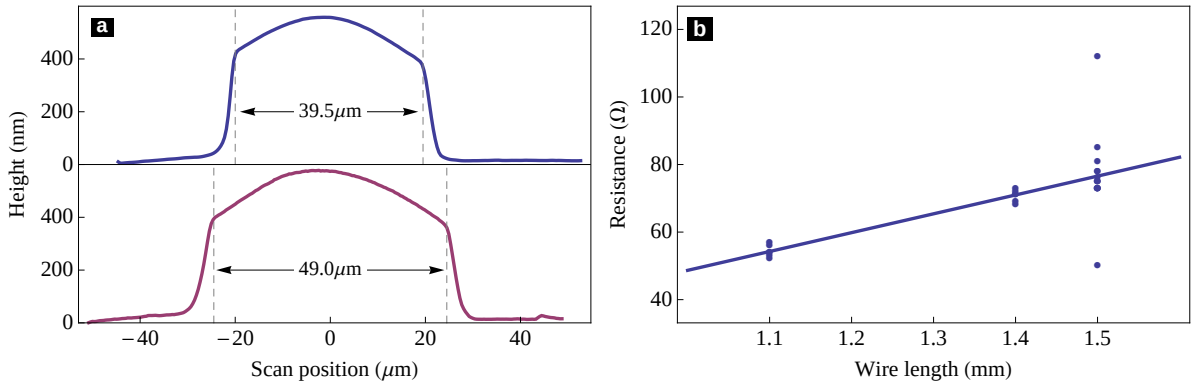

Figure 2.16: **(a)** AFM profile of wires made with 40 and 50 microns masks. **(b)** resistance of a number of 50 microns wires with different lengh on a single chip. The inferred resistivity of $1.3 \pm 0.3 \times 10^{-6}\,\Omega$m is close to the expected value.

We measured the resistance of nearly 30 50 μm-wide wires with length 1, 1.1 and 1.4 mm, all on the same sample. The results are presented in Figure 2.16(b). We estimate the resistance per unit length from a linear fit: $R_l = 56 \pm 11\,\Omega$/mm. Then, based on the wire cross-section inferred from the AFM measurements of Fig. 2.16(b), $A = 50 \times 0.45 = 23 \pm 2\,\mu\text{m}^2$, we obtain the material resistivity as $\rho = R_l \times A = 1.3 \pm 0.3 \times 10^{-6}\,\Omega$m, well within the typical resistivity range of NiCr, $1 - 5 \times 10^{-6}\,\Omega$m.

The wires were connected to a PCB via springy gold pins (Coda Systems Ltd), as illustrated in Fig. 2.17(b), which facilitates electric connections. We recorded the voltage-current characteristic for a number of heaters; a typical example is shown in Figure 2.17(a). The resistance remains constant even at high powers (several watts). Most wires can dissipate up to five watts of power before breaking down. Temperature of close to 1000 K can be achieved, as evidenced by the red glow we observed on several wires and in Fig. 2.17(b). Note that we only require a few tens of Kelvins to operate the heaters as phase shifters.

On the glass substrates, we measured the temperature by placing a thermocouple underneath the slide. A maximum temperature rise of 250°C was measured, after which the wires broke down. The pictures of broken wires shown on Figure 2.18 indicates that the breakdown is caused mainly by cracks in the glass slide, but we also observe a darkening of the wires that suggests oxidation. Much higher power/temperatures are achievable on waveguide chips, because the Si substrate underneath the thin ($\sim 50\,\mu$m) SiO_2 layer acts as a heat sink by virtue of its much higher thermal conductivity. Together with the aluminium heat sink on which the chip is mounted, this dissipates the heat very effectively. No structural damage to the sample was observed even after the wires broke down when fed 5 W of electrical power. Instead the wires turned completely black with a rugged/chipped aspect, suggesting that breakdown results from a fast and thorough oxidation. No noticeable change in the resistance is observed even at high power. This oxidation process comes as a surprise as NiCr is known to be extremely resistant to oxidation , thanks to the protective chromium oxide layer it naturally develops (127).

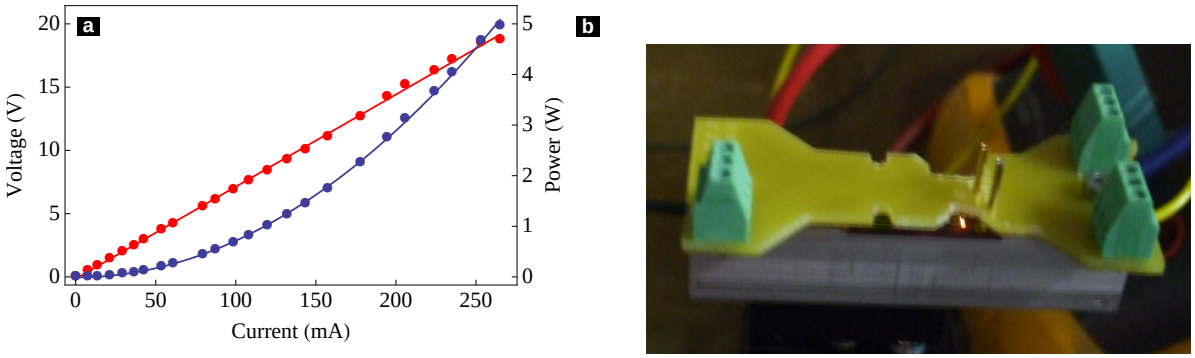

Figure 2.17: Crash test of a NiCr wire on a waveguide chip: ramp up the applied voltage until it blows up. **(a):** the $U(I)$ characteristic. The wire resistance is $dU/dI = 72\,\Omega$. Also shows is the dissipated power $P = UI$, which reaches 5 W before the current drops abruptly to zero. **(b):** Photograph of a waveguide chip mounted on the aluminium heat sink and connected to a PCB via springy gold pins. A wire can be seen glowing bright red, so the temperature must have reached about 1000 K

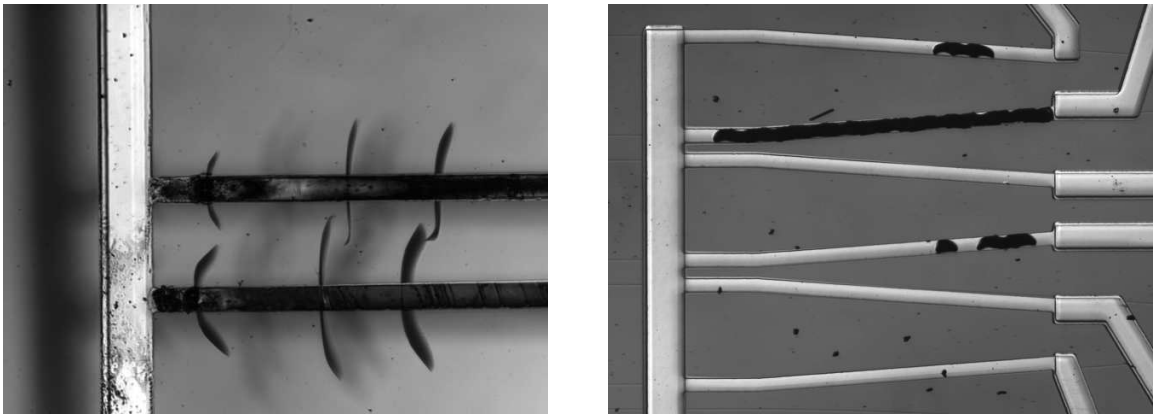

Figure 2.18: Burnt-out heaters. Left, an early experiment on a glass slide. The breakdown is a combination of oxidation and thermally induced strain and damage to the substrate The maximum temperature, measured with a thermocouple immediately underneath the glass slide, was 250°C. Right: on the latest chip, ORC8. The silica-on-silicon structure is much more efficient at dissipating temperature and substrate damage is not an issue, while the oxidation is more complete. There is large variation in the breakdown power; the best wires can dissipate close to 5 W.

Before we can measure the heaters' phase shift and bandwidth, we need to build cavities into the waveguides. This is the subject of the next chapter.

Chapter 3

Single cavities

In this short chapter we present the basic theory of lossy Fabry-Pérot resonators, a prerequisite for understanding the coupled cavity systems which will be the subject of the next chapter.

The phenomenon of interference played a crucial role in the debate over the nature of light that raged for several centuries. The observation of white light interferences (Newton rings) by Boyle and Hook dates back to the 17th century, but it was not before the dawn of the 19th century that the wave theory was put on a solid footing by Young's double-slit experiment. By 1820 the observation of Fresnel's (or Arago's) bright spot secured victory for the wave theory. Multi-beam interferences were studied much later still, as Charles Fabry and Alfred Pérot introduced their now ubiquitous apparatus in the last year of the 19th century. See Born & Wolf (128) for a concise yet comprehensive historical survey of these developments.

First we review the theory of lossless Fabry-Pérot resonators, which we then extend to include propagation loss. We limit our discussion to the case of on-axis plane waves. We then apply this theory to waveguide cavities, which allows us to measure the propagation loss, the fibre-waveguide mode-matching efficiency and the heaters' phase shift. We then move on to study, theoretically and experimentally, fibre- and waveguide-coupled plane-concave microcavities.

3.1 Theory of lossless Fabry-Pérot cavities

Fabry-Pérot interferences are extensively covered in many textbooks. On resonance, the confinement of light between two mirrors results in an increase of the intra-cavity field intensity and in destructive interferences between the field leaking out of the cavity and that being reflected at the input mirror, so that the cavity becomes transparent.

Consider plane waves with k-vector $k = k_0 n = 2\pi n/\lambda$ of the form $E(z) = e^{\pm ikz}$, propagating perpendicularly to the cavity mirrors, separated by a distance L, whose amplitude reflection coefficients are r_1 and r_2, as depicted in Figure 3.1. The refractive

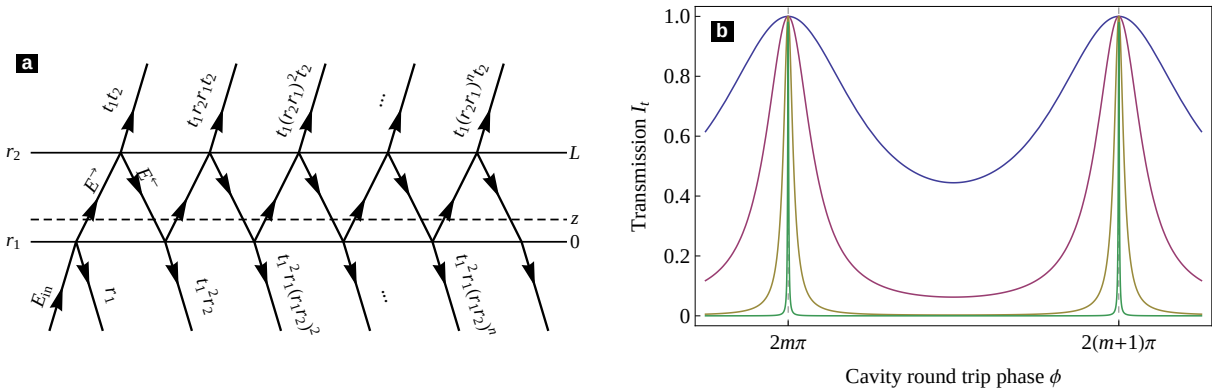

Figure 3.1: **(a)** Multiple reflections of a plane wave in a plane Fabry-Pérot resonator. **(a)** Transmission of such a cavity for identical mirrors of increasing reflectivities $R = \{20, 50, 90, 99\%\}$.

index of the medium is n and λ is the vacuum wavelength. We assume that $r_{1,2}$ are real, which implies lossless mirrors of zero thickness. The penetration length of the field inside the mirrors may be incorporated into the cavity length L. A wave of amplitude 1 impinging on the first mirror at $z = 0$ has an amplitude $t_1 = \sqrt{1 - r_1^2}$ immediately after the mirror. This wave will be reflected many times as it bounces back and forth between the mirrors, and the addition of these many reflected partial waves, as depicted in Figure 3.1(a), leads to the following expression for the forward- and backward-propagating travelling waves at position z inside the cavity:

$$E_{\text{cav}}^{\rightarrow}(z) = t_1 e^{ikz} \sum_{n=0}^{\infty} (r_1 r_2 e^{2ikL})^n = \frac{t_1 e^{ikz}}{1 - r_1 r_2 e^{2ikL}} \tag{3.1}$$

$$E_{\text{cav}}^{\leftarrow}(z) = -E_{\text{cav}}^{\rightarrow}(z) e^{ik(2L-z)} r_2 = -\frac{t_1 r_2 e^{ik(2L-z)}}{1 - r_1 r_2 e^{2ikL}} \tag{3.2}$$

These expressions are useful inasmuch as they also provides information about the intra-cavity standing wave. The reflected and transmitted fields are then given by

$$E_t = E_{\text{cav}}^{\rightarrow} e^{ik(L-z)} t_2 = \frac{t_1 t_2 e^{i\phi/2}}{1 - r_1 r_2 e^{i\phi}} \tag{3.3}$$

$$E_r = r_1 + E_{\text{cav}}^{\leftarrow} e^{ikz} t_1 = \frac{r_1 - r_2 e^{i\phi}}{1 - r_1 r_2 e^{i\phi}} \tag{3.4}$$

were we have introduced the round trip phase shift $\phi = 2kL$. Finally, the transmitted and reflected intensities are $I_{r,t} = |E_{r,t}|^2$ and

$$I_t = \frac{t_1^2 t_2^2}{1 + r_1^2 r_2^2 - 2 r_1 r_2 \cos\phi} \tag{3.5}$$

$$I_r = \frac{r_1^2 + r_2^2 - 2 r_1 r_2 \cos\phi}{1 + r_1^2 r_2^2 - 2 r_1 r_2 \cos\phi} \tag{3.6}$$

and we naturally have conservation of energy: $E_r + E_t = 1$, as long as the mirrors are lossless. It is convenient to introduce the *finesse coefficient* F by taking advantage of the trigonometric relation $\sin^2 \phi/2 = (1 - \cos\phi)/2$ as follows:

$$F = \frac{4r_1 r_2}{(1 - r_1 r_2)^2} \text{ such that } I_t = \frac{t_1^2 t_2^2}{(1 - r_1 r_2)^2} \cdot \frac{1}{1 + F \sin^2(\phi/2)} \tag{3.7}$$

This is the well known Airy function. It is plotted in Figure 3.1(b) for various reflectivities, and consists of equally spaced resonances which become sharper as the reflectivity increases. The resonances are spaced by the free spectral range $\Delta\phi = 2\pi$, most commonly expressed in the (physical) frequency domain as

$$\text{FSR} = \frac{c}{2nL}. \tag{3.8}$$

The FSR is nothing but the inverse of the photon round trip time.

The finesse coefficient F must not be mistaken for the *finesse* $\mathcal{F}$, which is the ratio between the full cavity linewidth $\delta\phi$ and free spectral range $\Delta\phi = 2\pi$. The linewidth $\delta\phi$ is defined as the full width at half-maximum (FWHM) of the intensity, ie

$$1 + F \sin^2(\pm\delta\phi/4) = 2 \Rightarrow \delta\phi = 4 \arcsin(F^{-1/2}). \tag{3.9}$$

If the finesse coefficient is high enough ($F \gtrsim 10$, which is almost always the case), the arcsin can be linearised so that $\delta\phi = 4/\sqrt{F}$. The finesse is then

$$\mathcal{F} = \frac{\Delta\phi}{\delta\phi} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2}. \tag{3.10}$$

The half-linewidth in the frequency domain, ie the cavity decay rate κ, is another important parameter. From $\delta\phi$ and ϕ we get

$$\kappa = \frac{c(1 - r_1 r_2)}{2nL\sqrt{r_1 r_2}} \simeq \frac{c\xi}{2nL} = \text{FSR} \times \xi \tag{3.11}$$

in *angular* frequency; we have introduced $\xi = 1 - r_1 r_2$, the loss of field amplitude in one cavity round trip due to leakage through the mirrors. In the last approximation the nature of κ as the round trip field loss per round trip time — i.e. a decay rate — is evident. It is also convenient to express κ as a function of the cavity finesse

$$\kappa = \frac{\pi c}{2nL\mathcal{F}} \tag{3.12}$$

The finesse is closely related to the mean number of round trips a photon would make in the cavity before leaking out, which is simply $\text{FSR}/\kappa = 1/\xi = (1 - r_1 r_2)^{-1}$. The resonant intra-cavity circulating intensity can become very large as a result of field

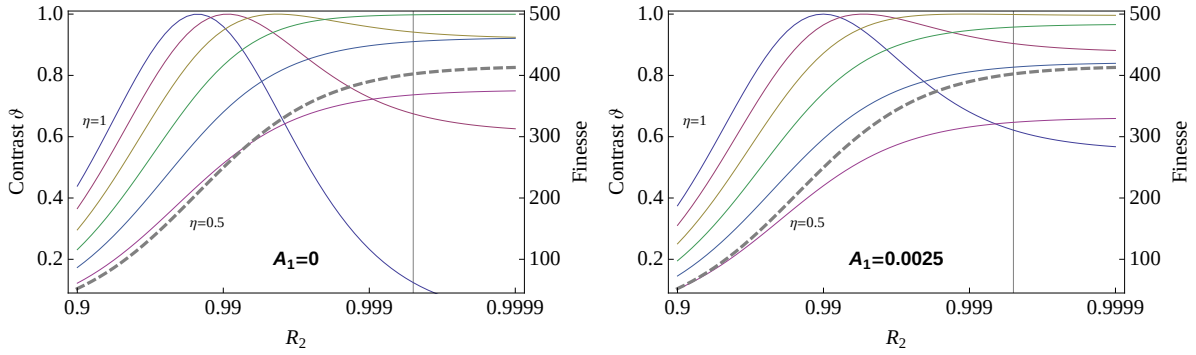

Figure 3.2: Contrast ϑ of the reflection fringes of a Fabry-Pérot cavity as a function of output mirror reflectivity R_2 and mode mismatch η. Left: lossless input mirror. Right: $A_1 = 0.25\%$ input mirror loss. The input mirror reflectivity is $R_1 = 98.5\%$. Note how the losses conspire to maintain a high contrast ($\vartheta > 0.5$) over a large range of parameters, even though the two mirror reflectivities are very different. The grey vertical line indicates the expected reflectivity of our spherical micro-mirrors ($R_2 \simeq 99.97\%$).

build-up. It is given by

$$I_{\text{circ}} = I_{\text{in}} \frac{T_1}{(1 - r_1 r_2)^2}. \tag{3.13}$$

Here $T_1 = |t_1|^2$, which is not necessarily equal to t_1^2. For example, with $R_1 = R_2 = 0.99$, $I_{\text{circ}} = 100 I_{\text{in}}$.

Another important parameter is the fringe contrast (or visibility) ϑ, which describes the modulation depth of the reflection spectrum. It is defined as

$$\vartheta = 1 - \frac{I_r^{\min}}{I_r^{\max}} = 1 - \left[\frac{(r_1 - r_2)(1 + r_1 r_2)}{(r_1 + r_2)(1 - r_1 r_2)}\right]^2 \simeq 1 - \left(\frac{r_1 - r_2}{1 - r_1 r_2}\right)^2. \tag{3.14}$$

Note that a somewhat more standard definition of the visibility would be $\vartheta = \frac{I_r^{\max} - I_r^{\min}}{I_r^{\max} + I_r^{\min}}$. Our definition (which we choose partly to be in line with our predecessors (129) leads to slightly simpler expressions. The last approximation in Eq. (3.14) is valid as soon as the reflectivities are sufficiently large, typically $R_1, R_2 > 0.9$. The contrast is plotted in Figure 3.2, where we have set $R_1 = r_1^2 = 0.985$ ($A_1 = 0, \eta = 1$ curve) and r_2 is the variable. The visibility reaches unity only when $r_2 = r_1$ (this is essentially an impedance-matching problem).

3.2 Theory of lossy Fabry-Pérot cavities

So far we have only considered a perfect resonator. However there are many sources of loss that can adversely affect the performance of the cavities. Let us consider three categories:

- **Mirror loss.** Everything that is not reflected off a *perfect* mirror gets transmitted: $R + T = 1$. But real mirror always suffer from a small loss A, such that $R +$

$T + A = 1$. In general this loss can be as small as 10^{-5} or better, and is usually negligible, especially for mirror of moderate reflectivities. We will neglect mirror loss throughout this thesis.

- **Propagation loss.** That is, a loss that originates from inside the cavity, independent from the mirrors. It can be caused by the natural absorption of the medium, by scattering off impurities, by the wall roughness in the case of a waveguide cavity, or even by the presence of ground state resonant atoms. We find that the most straightforward way of including this type of loss is to introduce an imaginary round trip phase Γ, so that in the previous equations ϕ is to be replaced by $\phi + i\Gamma$. Equivalently, one could include the loss in a complex refractive index, $n \to n + iq$. The conversion rules between the two, as well as between these and the loss in dB/cm, are given in Appendix C.

- **Mode mismatch.** This is not a cavity loss *per se* since it does not disturb the cavity (at least under some approximations to be detailed below), but a loss induced by the transverse mode mismatch between the cavity mode and the input fibre used to coupled light into the cavity.

3.2.1 Propagation loss

We start with the propagation loss. Substituting $\phi \to \phi + i\Gamma$ in $E_{r,t}$, we easily show that the transmission spectrum takes the form of an Airy function similar to Eq. (3.7), but where the finesse coefficient is now

$$F = \frac{4r_1 r_2 e^{-\Gamma}}{(1 - r_1 r_2 e^{-\Gamma})^2} \tag{3.15}$$

while the reflection contrast becomes

$$\vartheta = 1 - \left[\frac{(r_1 - r_2 e^{-\Gamma})(1 + r_1 r_2 e^{-\Gamma})}{(r_1 + r_2 e^{-\Gamma})(1 - r_1 r_2 e^{-\Gamma})}\right]^2 \tag{3.16}$$

These are exactly identical to Eq. (3.7) and (3.14) if we replace r_2 by $r_2 e^{-\Gamma}$. In other words, the intra-cavity loss is equivalent to a reduction in the reflectivity of the output mirror. Obviously we no longer have conservation of energy: $I_r + I_t < 1$. On the transmission side, the amplitude of the resonances is reduced by a factor

$$D = \frac{e^{-\Gamma}(1 - r_1 r_2)^2}{(e^{-\Gamma} - r_1 r_2)^2} \tag{3.17}$$

The significance of the imaginary round trip phase Γ we have introduced becomes clear if one consider the special case $r_1 = r_2 = 0$ (ie, by removing the mirrors to look

at propagation through a simple slab of absorbing material): the transmitted intensity is $e^{-\Gamma}$. Therefore $1 - e^{-\Gamma}$ is the single-pass fractional power loss.

Depending on whether the cavity was initially over-coupled ($r_1 > r_2$) or under-coupled ($r_1 < r_2$), the introduction of this propagation loss will improve or degrade the contrast, as shown in Figure 3.2. In the under-coupled case, unity contrast is reached when $r_2 e^{-\Gamma} = r_1$.

3.2.2 Mode mismatch

In this work we consider fibre-coupled waveguide resonators and waveguide-coupled plane-concave microcavities. In both cases, light impinging on the input mirror will excite the coherent superposition of resonant cavity modes $\{\Psi^n_{\text{cav}}\}$ that it best matches. Similarly, the cavity field that leaks through the mirror back to the input fibre/waveguide will excite a coherent superposition of guided modes $\{\Psi^m_{\text{in}}\}$. For our purpose it is sufficient to consider a single mode in the fibre, waveguide or microcavity so we drop the n, m indices. In general they are not perfectly mode-matched; ie the mode overlap is smaller than unity:

$$\eta = \iint \Psi^*_{\text{in}} \Psi_{\text{cav}} dxdy < 1 \tag{3.18}$$

and only a fraction $\eta^2 T_1$ of the intensity in the incoming mode will be coupled to the outgoing mode through the mirror (whose transmitivity is T_1^2). What happens then to the remaining fraction $(1-\eta^2)T_1$? It can either 1) be lost into cladding and leaky modes, with amplitude probability η_t — this we call the transmission limit, or 2) be rejected and reflected back to the original mode, with probability η_r (such that $\eta^2 + \eta_t^2 + \eta_r^2 = 1$); this is the reflection limit. Evaluating the relative contributions of η_t and η_r is an extremely complex task that requires computing all the cladding and leaky (non-guided) modes of the fibre/waveguide. However, the large number of these modes forms a quasi-continuum, so that we may assume $\eta_t \gg \eta_r$ and neglect case 2) altogether. This is different from the possibly more familiar case of a junction between two metallic microwave guides, where the field cannot leak out and therefore only the reflection limit is applicable.

The cavity reflection in the transmission limit is derived by simply substituting ηt_1 in lieu of t_1 in Eq. (3.1) and (3.4) (in bold below):

$$E^{\rightarrow}_{\text{cav}}(z) = \boldsymbol{\eta t_1} e^{ikz} \sum_{n=0}^{\infty} (r_1 r_2 e^{2ikL})^n = \frac{\eta t_1 e^{ikz}}{1 - r_1 r_2 e^{2ikL}} \tag{3.19}$$

$$\Rightarrow E_r = r_1 + E^{\leftarrow}_{\text{cav}} e^{ikz} \boldsymbol{\eta t_1} = \frac{r_1 - r_2 e^{i\phi}[1 - t_1^2(1-\eta^2)]}{1 - r_1 r_2 e^{i\phi}} \tag{3.20}$$

This means that a fraction η^2 of the intensity in the input mode goes into the cavity mode after being transmitted through the mirror, and symmetrically a fraction η^2 couples to the input mode after leaking through the input mirror. The remaining is lost forever.

The finesse is not affected, but the contrast will be reduced. It becomes

$$\vartheta = 1 - \left[\frac{(r_1 - (1-l)r_2)(1+r_1r_2)}{(r_1 + (1-l)r_2)(1-r_1r_2)}\right]^2 \tag{3.21}$$

with $l = t_1^2(1-\eta^2)$. The evolution of the contrast as a function of mode mismatch is shown in Figure 3.2. An interesting consequence of mode mismatch is to *increase* the visibility if the cavity is undercoupled ($R_1 < R_2$), which is the case of our microcavities (since we want photons to exit preferentially into the coupling fibre).

3.2.3 Mirror transmission loss

So far we have assumed perfect mirrors, that is we have substituted $\sqrt{1-r_1^2}$ for t_1 whenever the transmission appeared (and similarly for t_2). If the input mirror has a power loss A_1, this becomes $\sqrt{1-r^2-A_1}$ and the contrast becomes

$$\vartheta = 1 - \left[\frac{(r_1 - (1-A_1)r_2)(1+r_1r_2)}{(r_1 + (1-A_1)r_2)(1-r_1r_2)}\right]^2 \tag{3.22}$$

We see that it is formally equivalent to the mode matching loss in Eq. (3.21), with $A_1 \leftrightarrow l = t_1^2(1-\eta^2)$.

Figure 3.2 illustrates the interplay between mirror loss, mode matching and mirror mismatch. It shows the fringe visibility as a function of the output mirror reflectivity R_2, for $0.5 < \eta < 1$ and in the absence/presence of mirror loss A_1. The input mirror reflectivity is set to $R_1 = 98.5\%$. We see that without any loss, the visibility in the neighbourhood of $R_2 = 99.9\%$ (these values would be typical of our microcavities) is very poor, but that the introduction of some loss, either through mode mismatch or mirror loss (or both), will compensate for the mirror mismatch. In fact it is quite remarkable that in the region of interest ($R_2 \geq 0.99$) the visibility does not go below 50% over a very large range of parameters.

3.2.4 Experimental measurement of loss and reflectivity

It is possible to solve equations Eq. (3.15) and (3.17) or (3.16) to express the reflectivity and propagation loss as a function of the finesse coefficient and on-resonance transmission drop. This is interesting because the two latter quantities are experimentally much more accessible than the former, and it will come handy in chapter 5. We get

$$R = 1 + \frac{2D}{F}(1-\alpha) \tag{3.23}$$

$$\gamma = \frac{1}{F^2}\left(2D(F+2) - 4\alpha + F(F+2-2\alpha) + 2\sqrt{(F+1)(8D^2 + 8D(F-\alpha+F(F-4\alpha)}\right) \tag{3.24}$$

with $\alpha = \sqrt{D(D+F)}$. A modestly simpler form can be reached through the approximation $F >> D$ (since $F >> 1$ and $D < 1$). Inclusion of the mirror loss A is possible but yields complicated expressions that do not need to be reproduced here.

3.2.5 Cooperativity

Although of limited practical use, it is an interesting exercise to calculate the cooperativity for an atom maximally coupled to a waveguide cavity (i.e. the atom sits in the centre of the waveguide mode). Remember that it is given by Eq. (1.17), and using Eq. (3.12) and (1.7), we have simply

$$C = \frac{g^2}{2\kappa\gamma} = \frac{1}{2\gamma} \cdot \frac{\mu^2\omega}{2\hbar\varepsilon_0 V} \cdot \frac{2L\mathcal{F}}{\pi c} = \frac{1}{2\gamma} \cdot \frac{\mu^2\omega}{2\hbar\varepsilon_0\pi c} \cdot \frac{2}{\pi c} \cdot \frac{\mathcal{F}}{A} \tag{3.25}$$

This expression highlights an interesting fact: that the cooperativity is to some extent independent of the cavity length L, because although the mode volume $V = A \times L$ gets larger with L the resulting decrease in g^2 is compensated by the corresponding reduction in the cavity decay rate κ. In this case, L cancels out completely because the mode area is constant, which is not the case in a free space cavity (see below).

3.3 Dielectric mirrors

All mirrors used throughout this work are standard thin film dielectric coatings, deposited on the carefully polished end facets of waveguide chips (on which waveguides had previously been written), and on the spherical micromirror arrays. The thin films were deposited by OIB GmBH in Jena, Germany, and consist of alternate layers of MgO (refractive index $n_L = 1.75$) and TiO_2 ($n_H = 2.65$).

The high reflectivity of these distributed Bragg reflectors (DBRs) originates from the constructive interference of plane waves in a stack of low and high-index quarter-wave layers. They are usually described by one or another variant of the traditional transfer matrix method (we like the k-vector formalism developed in (130)), in which each interface between two dielectrics and each dielectric layer is represented by a 2×2 matrix relating the forward and backward-propagating fields on each side of the element. Arbitrary planar structures can be described simply by multiplying the appropriate matrices.

The matrices M_{prop} describing propagation in a dielectric layer of thickness d and refractive index n_1, and M_{interf} describing the interface between two dielectrics of refractive indices n_1 and n_2, are

$$M_{\text{prop}} = \begin{pmatrix} e^{-idk_0 n_1} & 0 \\ 0 & e^{idk_0 n_1} \end{pmatrix} \text{ and } M_{\text{interf}} = \frac{1}{t_{12}} \begin{pmatrix} 1 & -r_{21} \\ r_{12} & t_{12}t_{21} - r_{12}r_{21} \end{pmatrix}. \tag{3.26}$$

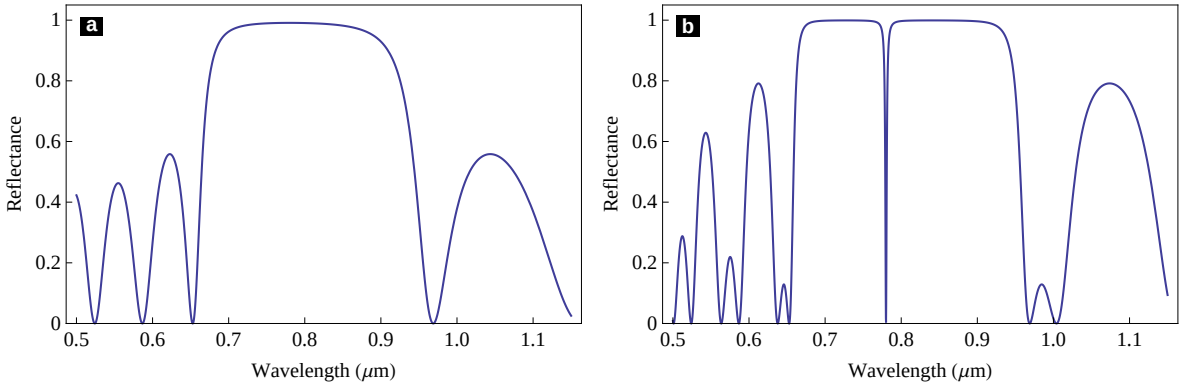

Figure 3.3: Left: Reflectance spectrum of a DBR mirror made up of 5 bilayers. The design wavelength is 780 nm, the refractive indices are 1.75 and 1.65, both substrates are air. Right: two such mirrors have been stack together with a $\lambda/2$ gap, forming a cavity.

r_{ij} and t_{ij} are the usual Fresnel reflection and transmission coefficients

$$r_{ij} = \frac{n_i - n_j}{n_i + n_j} \text{ and } t_{ij} = \frac{2n_i}{n_i + n_j} \tag{3.27}$$

We limited ourselves to normal incidence, which is sufficient for our purpose. The method can be straightforwardly generalised to s and p-polarised waves at other incidence angles.

As an example, we plot in Figure 3.3 the calculated reflection spectrum of two different structures. (a) is a simple mirror consisting of five pairs of high/low index quarter-wave layers at 780 nm. The reflectivity at the design wavelength is 98.5%, and the bandwidth is close to ~ 200 nm. This is very large compared with the Bragg gratings we discuss in Chapter 5, thanks to the large index contrast. The side lobes originate from higher-order interferences. In (b) we concatenated two of these mirrors, with a half-wave gap. The resulting cavity makes the structure transparent at the design wavelength.

3.4 Waveguide cavities

3.4.1 Optimising the mirror reflectivities

Initial devices and experiments

For a first experimental run we coated mirrors with reflectivities 99% and 99.5% [1] on several chips, including "CIP" (fabricated by deposition and etching, length 32 mm) and "Jena" (femtosecond-written, length 20 mm). These values were chosen according to the theoretical investigation presented in Chapter 1.5.3, without considering the propagation losses which were not yet known.

Normalised reflection spectra for these two chips are presented in Figure 3.4. The

[1] This should have been 98% but for a manufacturing error outside of our control.

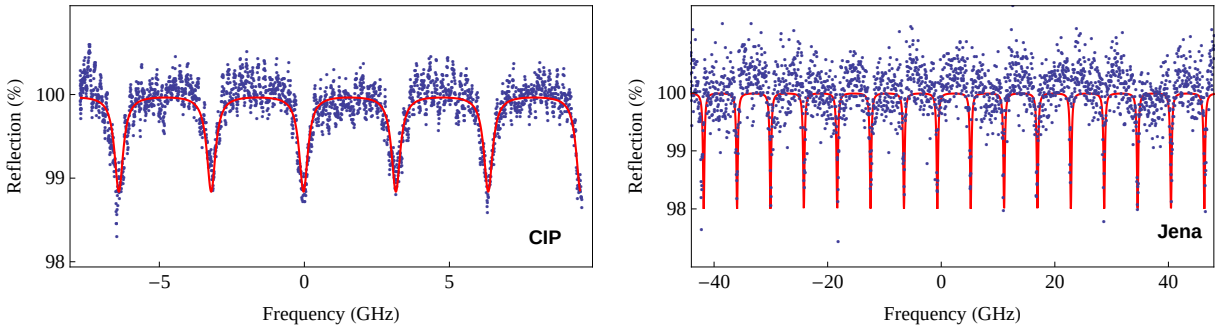

Figure 3.4: Typical normalised reflection spectrum of waveguide cavities on the "CIP" and "Jena" chips. The contrast is very low because both mirrors have very similar reflectivities (99 and 99.5%) together with a significant propagation loss. CIP chip: the finesse is 10 ± 1 and the contrast is $1.2 \pm 0.1\%$, from which we infer propagation loss of 0.5 dB/cm and a mode-matching efficiency $\eta_w^2 = 42\%$. Jena chip: $\mathcal{F} = 24 \pm 1$ and $\vartheta = 2.0 \pm 0.5\%$, hence the loss is 0.28 ± 0.02 dB/cm and $\eta_w^2 = 0.29 \pm 0.07\%$. As shown in the previous chapter, these waveguides are much larger and therefore the mode matching is rather poor.

red lines are Airy function fits to the data. They give a free spectral ranges of 3.2 and 5.1 GHz respectively, in accordance with the waveguide lengths and with an effective index $n_{\text{eff}} = 1.46$. From the fits we also extract the finesses $\mathcal{F} = 10 \pm 1$ and 24 ± 1. In the absence of losses, we would expect the finesse to be $\mathcal{F} = 417$. With Eq. (3.15) we use the difference to infer the propagation loss as 0.6 and 0.3 dB/cm for the CIP and Jena chips respectively. 0.3 dB/cm is typical for femto-second written waveguides (122; 131). On the other hand 0.6 dB/cm is much higher than the upper limit of 0.1 dB/cm claimed by the manufacturer, CIP Technologies (132). But there are reasons to doubt the accuracy of this claim. In particular it was noticed in the CCM waveguide chip experiment (106), which used an almost identical chip, that roughly half of the expected waveguide transmission was unaccounted for, even after inclusion of all the other known sources of loss. As it turns out, the missing 50% may well be attributed to this 0.6 dB/cm factor, which translates to a 40% single pass loss.

The propagation loss on UV-written ORC chips was measured in a similar way, but with integrated Bragg gratings instead of end-facet mirrors. See Chapter 5 for more details. It was found to be 0.9 ± 0.3 dB/cm. This result was confirmed by independent measurements carried out at the ORC, using the method of Ref. (115).

The visibility of the spectra in Figure 3.4 is very poor: $\vartheta = 1.2 \pm 0.1\%$ and $2.0 \pm 0.5\%$ respectively. This is because the large loss effectively reduces the reflectivity of the second mirror and, as we have seen, mirror mismatch results in low contrast. This makes measurements somewhat difficult, in part because residual low-finesse interference fringes from the coupling fibre are on a similar level (it is the main source of noise observed in these data).

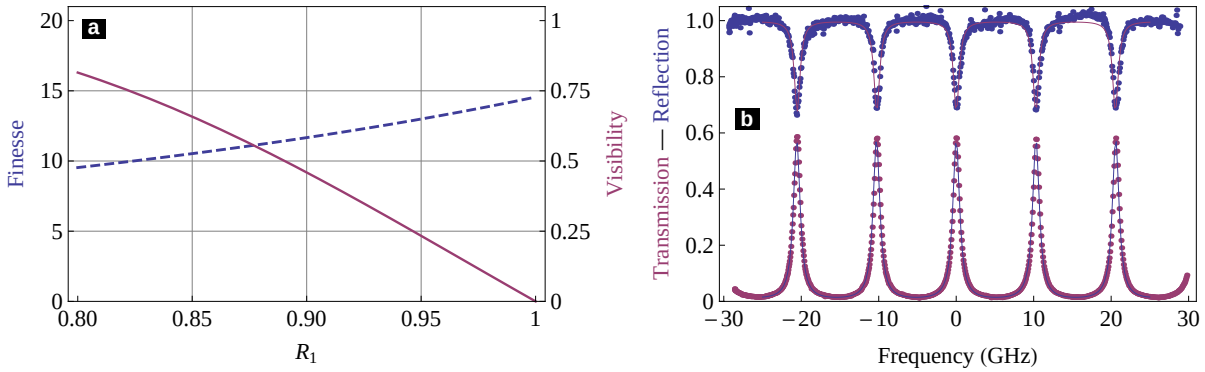

Figure 3.5: Optimised input mirror. **(a)** Calculated finesse (dashed blue line) and fringe visibility (solid purple line) of a cavity with 0.9 dB/cm propagation loss and output mirror $R_2 = 0.99\%$, showing that it is sensible to decrease R_1 from about 99% to about 90% to increase the visibility without compromising the finesse too much. **(b)** Reflection (top; normalised signal) and transmission (bottom; arbitrary units) spectra of waveguide cavity ORC7-A1. Fitting these with a model including propagation loss and mode mismatch between the waveguide and the coupling fibre, we can estimate the first as 0.9 ± 0.1 dB/cm and the second as $\eta_w^2 = 0.48 \pm 0.01$.

Optimised mirrors

In order to improve the visibility, we made a few changes for the next generation of chips (UV-written, codename "ORC7"). First we wrote waveguides along the short chip dimensions (1 cm) to minimise the loss. Secondly, we decreased the reflectivity of the input mirror R_1 (while keeping the output mirror $R_2 = 99\%$) to reduce the mirror mismatch. Figure 3.5(a) shows how the finesse (dashed blue) and the visibility (purple curve) are expected to change as a function of R_1, given propagation loss of 0.9 dB/cm. As we can see, with the initial $R_1 \simeq 99\%$, the visibility is at most a few per cent, and the finesse is 14. By decreasing R_1, the visibility can be improved substantially without compromising the finesse too much, as long as the overall damping is dominated by the propagation loss. We settled for $R_1 = 90\%$, at which point the visibility is about 40%, while the decrease in finesse ($\mathcal{F} = 12$) is negligible.

With this parameters, we wrote straight waveguides and deposited mirrors on chip ORC7. Figure 3.5(b) shows a typical transmission and normalised reflection spectrum. The easiest way of extracting the required information from these is to normalise the spectra to remove the laser intensity variation during a scan, so that the off-resonance reflection is unity. We then fit a model based on the analysis in the previous section, incorporating both propagation loss and mode mismatch between the waveguide cavity and the coupling fibre.

From a number of these fits we extract the finesse $\mathcal{F} \sim 10 \pm 2$ and the propagation loss 1.0 ± 0.2 dB/cm, in line with the other experiments we already mentioned. The mode-matching efficiency is estimated to be $\eta_w^2 = 0.48 \pm 0.01$, a very disappointing result when one remembers that we expected $\eta_w^2 \gtrsim 0.9$. This discrepancy might be caused by

the dielectric mirror, but then we would think that the mirror reflectivity would also be affected — which does not seem to be the case. This is therefore a bit of a mystery, and one which did not have the leisure to explore in any greater details.

3.4.2 Tuning the cavity length with phase shifters

We introduced the thermo-optic phase shifters in the previous chapter, but as long as we did not have waveguide cavities, it was impossible to actually prove that they acted as such. Here we demonstrate phase shifts of several π and investigate the response time.

We deposited NiCr heaters on top of several waveguide cavities on chip ORC7. As a power supply we used a three-channels piezo driver (Thorlabs BPC203, maximum output 75 V, 500 mA). This lets us control several heaters simultaneously; the driver can be controlled either manually, remotely from the computer or with a function generator.

Phase shift

Figure 3.6(a) shows the reflection spectrum of a waveguide cavity as the power dissipated in the heater increases. The resistance of the wire was $R = 107\,\Omega$. The spectrum shifts to lower frequencies as the control voltage V, indicated on the vertical axis, increases. The shift reaches one full FSR (i.e. a single pass phase shift of π) at $V \sim 5.5\,\mathrm{V}$, or a dissipated power $P = V^2/R \sim 0.28\,\mathrm{W}$.

In Figure 3.6(b) we plot the frequency shift (also expressed as a phase shift on the right axis) as a function of control voltage (bottom axis) or power (top axis), for two different heaters. The blue one is 1.1 mm long while the red one is 1.5 mm. We see that the phase (frequency) shift is quadratic in the voltage, but linear in the power. Indeed it should be proportional to the temperature change, which is itself proportional to the dissipated power. The maximum phase shift here is about 3π or three full FSRs.

For a given voltage/power, the phase shift is larger for the longer heater (red). However the slope $\delta\phi/dP$ depends on other factors than the length, such as the wire resistance (which can vary by 20% at constant length) and the relative position of the heater and the waveguide. The latter is only 5 μm wide whereas the heaters width is 40 or 50 μm, and they are manually positioned on top of the waveguide with an accuracy no better than 10 to 20 μm. In general $\delta\phi/dP$ is 3 to 5π/W. Because of these variations it is necessary to calibrate each heater indpendently.

Response time

We measured the phase shifters' bandwidth in two ways. First we turned off the laser frequency scan and sat the laser on the side of a resonance fringe, and observed the reflection signal as a small square voltage modulation is applied to the heater. The resonance frequency shift results in a small change in the reflection amplitude (Fig. 3.7,

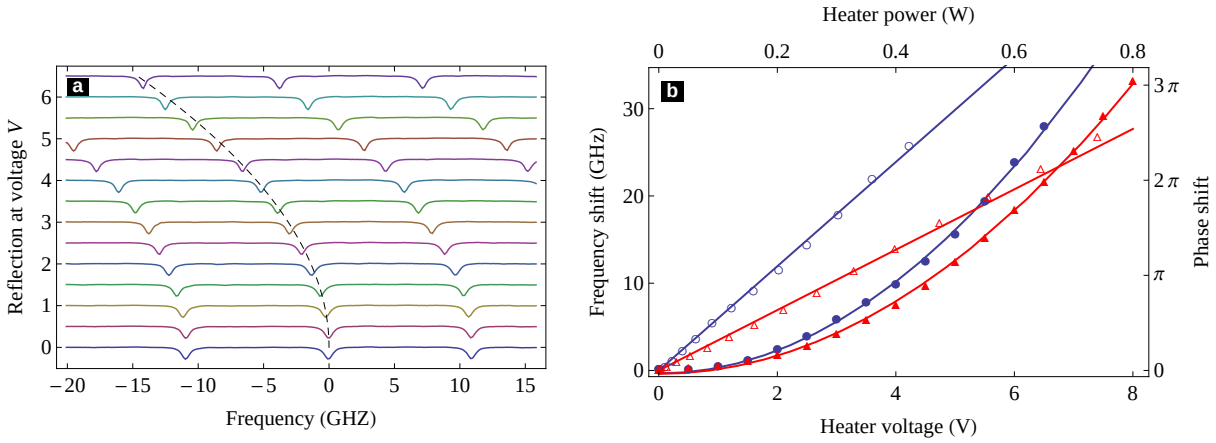

Figure 3.6: **(a)** Frequency shift of a waveguide cavity reflection spectrum as a function of heater voltage (indicated on the vertical axis.) **(b)** Frequency shift Δf and phase shift ($\Delta\phi = \pi\Delta f/FSR$) as a function of heater voltage V and power V^2/R. The red curves correspond to the spectra on the left (wire length 1.5 mm, resistance $R = 110\,\Omega$. The blue curves are for a different heater (1.1 mm, 75 Ω).

left), which is then fitted with a simple exponential $\exp(-t/\tau)$. We measured a 60% rise time $\tau = 0.66 \pm 0.05$ ms.

In a second experiment, the laser frequency was also adjusted to sit on the side of a fringe, but this time the heater voltage modulation was sinusoidal and we recorded the resulting modulation amplitude (gain) as a function of modulation frequency (green waveforms and blue dots). We then treat the heater as a simple first order, low-pass RC circuit with a ("voltage") gain function

$$G(f) \propto \frac{1}{\sqrt{1 + (2\pi f\tau)^2}}. \tag{3.28}$$

This fit gives a more precise response time $\tau = 0.56 \pm 0.01$ ms that is compatible with the previous result. The frequency at which the modulation amplitude has dropped by half [2] is $f_{1/2} = \sqrt{3}/(2\pi\tau) = 492 \pm 10$ Hz. We note that this compares very favourably with other silica-on-silicon thermo-optic phase shifters that have been reported in the literature, and whose response time typically lies in a 1-10 ms range.

The bandwidth of our heaters being of little importance at the moment, we did not attempt a systematic study of many devices.

[2] In an actual electronic filter we would consider the 3 dB frequency as the point where the *power* has dropped by half, or $G^2(f_{3\text{dB}}) = G^2(0)/2 \Rightarrow f_{3\text{dB}} = 1/2\pi\tau$. But the "power" is ill-defined in our case.

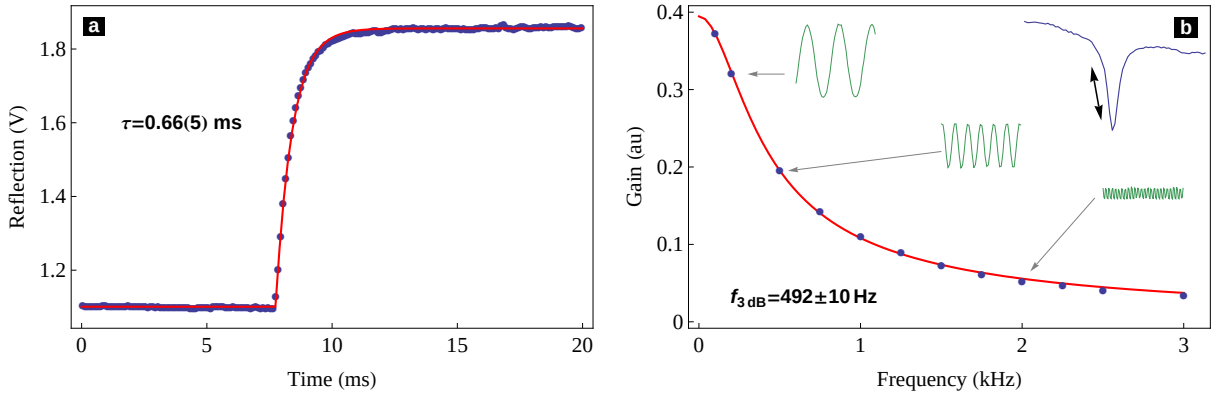

Figure 3.7: Measurement of a phase shifters' response time. **(a):** the laser sits on the side of a resonance fringe, and a small square modulation is applied to the heater. **(b):** the modulation is now sinusoidal, and the signal's modulation amplitude is recorded as a function of modulation frequency. It is then fitted with a simple first-order, low-pass filter gain function. The half-maximum cut-off frequency is 492(10) Hz.

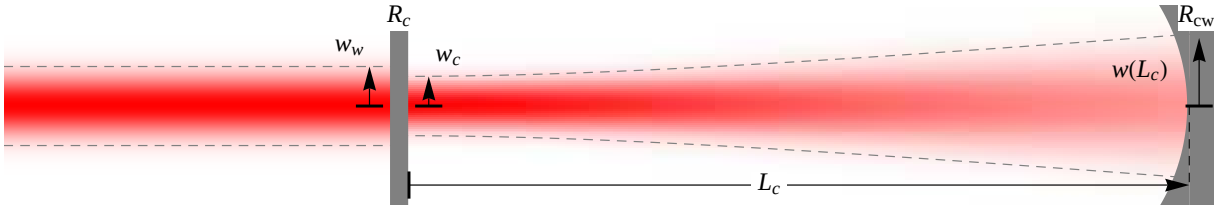

Figure 3.8: Schematics of a waveguide coupled to a microcavity.

3.5 Waveguide- and fibre-coupled microcavities

3.5.1 Plane-concave microcavities

Having dealt with the waveguide resonators, we now turn our attention to the next type of cavity in our system, the free space microcavities where the atoms will eventually sit. They are simple plane-convex Fabry-Pérot resonators as depicted in Figure 3.8. In the following we assume that the reader is familiar with the optics of Gaussian beams, which is summarised in Appendix A.

Here the plane mirror is a multilayer coating on the end facet of the waveguide chip, or on the tip of an optical fibre, with reflectivity R_{cw}. The cavity is closed by a spherical mirror on the micromirror array chip, which has a radius of curvature ρ_{c} and is covered with a high-reflectivity coating R_{c}. The distance between the two mirrors is L_{c}. Ignoring all higher-order modes, the cavity mode is a paraxial Gaussian beam with a waist of radius w_{c} located at the plane mirror ($z = 0$). The beam radius $w(z)$ and the radius of curvature $\rho(z)$ at position z along the cavity are:

$$w(z) = w_{\mathrm{c}}\sqrt{1 + (z/z_R)^2} \tag{3.29}$$

$$\rho(z) = z + z_R^2/z \tag{3.30}$$

$z_R = \pi w_c^2/\lambda$ is the Rayleigh range: $w(z_R) = w_c\sqrt{2}$.

The cavity is stable only if $0 < L_c < \rho_c$. The waist w_c is a function of the cavity length: by matching the beam curvature at L_c to that of the spherical mirror, one can solve $\rho(L_c) = \rho_c$ to get

$$w_c(L_c) = \sqrt{\lambda/\pi}\sqrt[4]{L_c(\rho_c - L_c)}. \tag{3.31}$$

Figure 3.9(a) shows the evolution of w_c as a function of cavity length for a mirror with radius of curvature $\rho_c = 110\,\mu$m. It reaches a maximum at $L_c = \rho_c/2$ and goes to zero at $L_c \to 0$ and $L_c \to \rho_c$, where the mode diverges abruptly (the mode radius at the mirror is shown in blue). The maximum value of w_c is $w_c^{\max} = \sqrt{\rho c\lambda/2\pi}$.

In our case, with $\rho_c = 110\,\mu$m, $w_c^{\max} = 3.7\,\mu$m. This is typically larger than the fibre or waveguide waist $w_w \sim 2.5\,\mu$m, so there will be two cavity lengths for which the cavity and the waveguide are mode-matched ($w_c = w_w$), as illustrated in Fig. 3.9(a). This is a favourable situation since it minimizes the losses. The optimum cavity lengths are

$$L_c^{\pm} = \frac{\rho_c}{2} \pm \sqrt{\frac{\rho_c^2}{4} - z_{R,w}^2} \tag{3.32}$$

$z_{R,w} = \pi w_w^2/\lambda$ is the Rayleigh range of the waveguide mode. In Figure 3.9(a) we also plot the mode-matching efficiency η^2 from Eq. (A.7). $\eta^2 = 1$ when $L_c = L_c^{\pm}$. It is interesting to see that even in the worst case $L_c = \rho_c/2$, the mode-matching efficiency does not fall much below 85%: even though it should be advantageous to work close to $L_c^{\pm}$, this is not a critical parameter.

We refer to L_c^- and L_c^+ as the short and long cavity regime, respectively. The large cavity regime is often easier to achieve and provides access to the cavity from the sides, but the larger mode radius on the mirror can be detrimental to the finesse if the mirror is too small or not perfectly spherical over a wide enough area. On the other hand, the short cavity regime offers a higher cooperativity, by virtue of its smaller mode volume, as seen in Fig. 3.9(b) and calculated below. However, we see that L_c^- lies in the grey-shaded area, which means that the cavity is shorter than the depth of the mirror. This region is not accessible because the diameter of the fibre is larger than the diameter of the mirror. Table 3.1 summarises the parameters of the four different micromirror arrays at out disposal. None of them allows access to L_c^-.

We should point out that microcavities similar to ours have been developed elsewhere, but in most other realisations the curved mirror is on the fibre side. This can be done by etching the fibre tip with a CO_2 laser (133; 134) or by transferring and gluing a pre-curved mirror (135). However in this geometry there is a fundamental mode mismatch between the cavity and fibre modes, which probably explains why in these works the collection efficiency does not exceed 10%. In contrast, our waveguide-coupled geometry is intrinsically mode-matched.

Sample	ρ_c (μm)	d (μm)	s (μm)	L_c^- (μm)
1	100	100	13.4	6.8
2	110	113	7.7	6.1
3	210	86	8.7	3.1
4	310	145	8.6	2.1

Table 3.1: Parameters of the four micromirror arrays available: radius of curvature ρ_c, aperture diameter d, depth $s = \rho_c - \sqrt{\rho_c^2 - d^2}$, short cavity mode-matching length L_c^-.

Cooperativity

We can also calculate the cooperativity. We consider two different cases: a single atom, located close to the microcavity waist, and an atomic ensemble distributed over the whole cavity mode. In the first case, the mode volume is still $A \times L_c$ (since the mode volume is defined by the cross-section at the position of the emitter), but now $A = \pi w_c^2/2$ is a function of cavity length. The cooperativity is then, using Eqs. (3.12) and (1.7)

$$C = \frac{g^2}{2\kappa\gamma} = \frac{2\mathcal{F}\mu^2}{\gamma\lambda^2\varepsilon_0\hbar}\frac{1}{\sqrt{L_c(\rho_c - L_c)}}. \tag{3.33}$$

In the second case, with atoms distributed over the entire cavity, we have to integrate over the whole optical mode. The mode volume is found to be

$$V = \int_0^L \frac{\pi w(z)^2}{2} dz = \frac{\lambda\sqrt{L_c^3(\rho_c - L_c)}}{2}\left(1 + \frac{L_c^2}{3z_R^2}\right). \tag{3.34}$$

and the cooperativity *per atom* is

$$C = \frac{g^2}{2\kappa\gamma} = \frac{6\mathcal{F}\mu^2}{\gamma\lambda^2\varepsilon_0\hbar}\frac{\sqrt{\rho_c - L_c}}{\sqrt{L_c}(3\rho_c - 2L_c)}. \tag{3.35}$$

The cooperativity for typical parameters of our microcavities is plotted in Figure 3.9, with a finesse $\mathcal{F} = 1000$ (corresponding to $R_{1,2} \simeq 0.997$) and a mirror curvature $\rho_c = 110\,\mu$m. In the case of an atomic ensemble (blue curve), it reaches about 5 in the short cavity regime, and drops to about 1-2 for a long cavity. $C \simeq 1$ is enough to observe some Purcell enhancement. Note that in the vicinity of $L = \rho_c/2$, $C \sim 4$ and does not vary significantly over quite a large range. If we had assume a non-expanding Gaussian beam, Eq. 3.25 would have given us a similar value — so this approximation can be useful. For a single atom (purple curve), the optimum would be close to the stability limit in the long cavity regime, where $w_c \to 0$ and the cooperativity diverges. C will be limited in practice by the extreme divergence of the cavity mode for small w_c, which will induce a large loss at the spherical mirror.

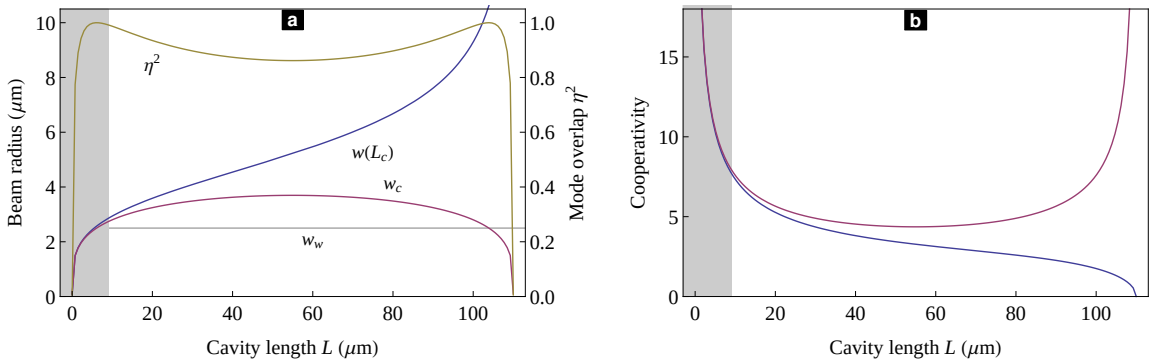

Figure 3.9: Cavity parameters as a function of cavity length, for a micromirror radius of curvature $\rho_c = 110\,\mu$m and a finesse $\mathcal{F} = 1000$. **(a):** mode waist w_c, beam radius at spherical mirror $w(L_c)$ and mode overlap η^2 between the cavity and the waveguide, whose mode size w_w is also indicated. **(b):** Cooperativity C as a function of cavity length, for a dilute atom cloud (blue) and for a single atom close to the plane mirror. The micromirror depth is 7.7 μm, and so the shaded area is experimentally inaccessible.

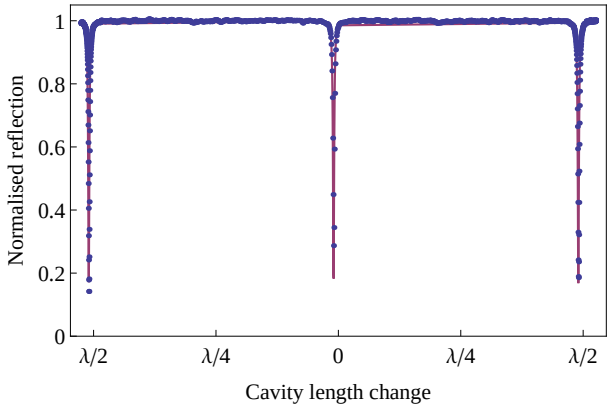

Figure 3.10: Reflection spectrum of a fibre-coupled microcavity. The fibred mirror reflectivity is 98.5% while that of the spherical micromirror is around 99.9%. The finesse is $\mathcal{F} = 230$, close to the maximum theoretical finesse of 300.

3.5.2 Microcavity length scan

We built and characterised plane-convex microcavities that were coupled either to a single fibre (Nufern HP-780) or to waveguide chips. In the latter case, only one mirror was deposited on the chip to form the plane mirror; the other end of the waveguide was butt-coupled to a fibre. In order to align and tune the length of the microcavity, the spherical mirror was mounted on a three-axis piezo flexure stage (Thorlabs NanoMax300). Because a cavity of length 20-100 μm has a free spectral range of a few THz, the finesse cannot be measured by scanning the laser frequency, as we did with the waveguide resonators. Instead we have to scan the cavity length, typically by driving the piezo sinusoidally at 50-100 Hz. A typical spectrum is shown in Figure 3.10.

A major problem with these NanoMax stages is the existence of low-frequency mechanical resonances that appear to be always excited — even when the piezo is not being driven. This is graphically by Fig. 3.11, which shows oscilloscope traces over a single cav-

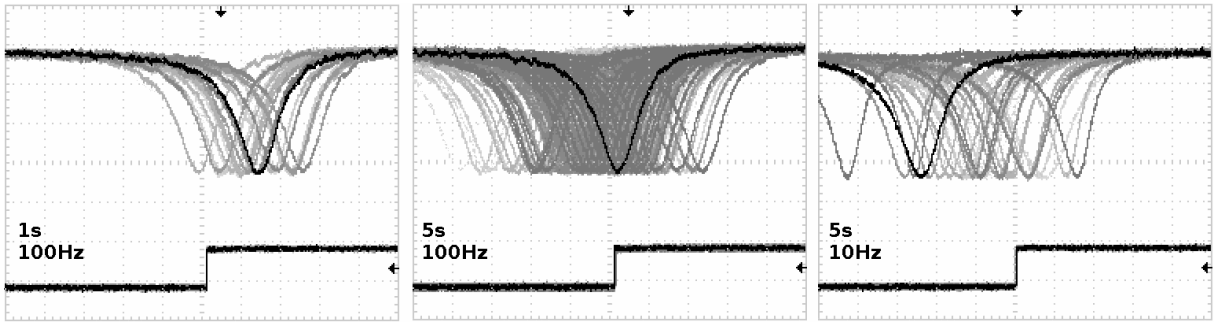

Figure 3.11: Oscilloscope traces over a microcavity fringes before the introduction of damping. These traces are zoomed in the linear region of the piezo sinusoidal modulation, and the time axis spans about 12 nm. The "persistence" function was set to 1, 5 and 5 s respectively, while the piezo modulation frequency was 100, 100 and 10 Hz. As we can see, the amplitude of the mechanical vibration is of the order of 5-6 nm, i.e. much more than the linewidth. The lower trace is the trigger signal.

ity fringe. The cavity in this case was waveguide-coupled. The amplitude of vibration is about 5-6 nm, whereas the linewidth is only about 1 nm. While mechanical resonances are expected in this type of flexure stages (as mentioned in the manufacturer's specification), it is surprising that they oscillate spontaneously. Attempts at isolating the stages from other sources of vibration and noise (acoustic box, floating the optical table, disconnecting the piezos, turning off nearby pumps and fans...) failed to produce any improvement.

In order to ascertain the origin of this vibration, we acquired a complete mechanical spectrum by driving the stage with a sine modulation of varying frequency, adjusting the drive amplitude so as to observe exactly one microcavity FSR over one modulation period. The result is plotted in Figure 3.12(b, inset), and we compare it with a simple mechanical harmonic oscillator model, whose amplitude spectrum is (136)

$$A(\omega) \propto \sqrt{\frac{1}{\gamma^2\omega^2 + (K - M\omega^2)^2}}. \tag{3.36}$$

k_B is the Boltzmann constant, $T = 300\,\mathrm{K}$ is the temperature and M, K and γ are the mass, spring constant and damping coefficient of the oscillator, respectively. The resonant frequency ω_0 and the quality factor Q are then given by $K = M\omega_0^2$ and $\gamma = M\omega_0/Q$. The fit gives $M = 65\,\mathrm{g}$, $K = 7900\,\mathrm{g s^{-2}}$, $\gamma = 0.11\,\mathrm{g s^{-1}}$, $\omega_0 = 350\,\mathrm{Hz}$ and $Q = 200$. While is is difficult to evaluate the mass of the oscillating part of the stage, 65 g certainly sounds reasonable, and the resonant frequency is compatible with that given in Thorlabs' specifications (375 Hz without a load; in this case the load consisted only of the micromirror array and a small mount, weighing at most a few tens of grams).

We then turned off the piezo driver, and disconnected the piezo altogether to eliminate the electrical noise. We scanned the laser frequency around a cavity resonance at 2 kHz. We recorded many spectra like the one shown in Fig. 3.12(a) and monitored the position of the fringe, $f_0(t)$. By Fourier-transforming $f_0(t)$ we obtain the top spectrum in Fig. 3.12(b).

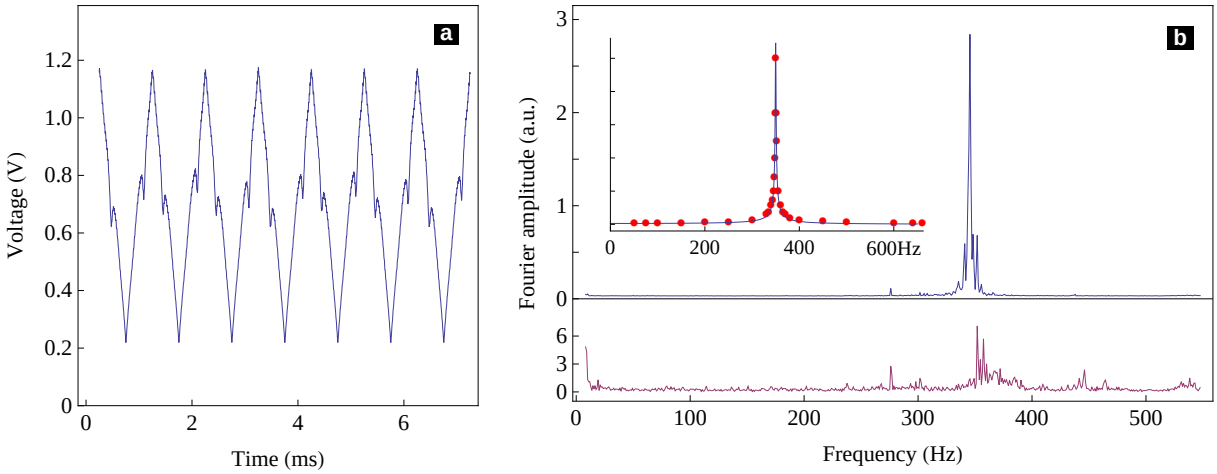

Figure 3.12: **(a)**: Reflection from the microcavity as the laser frequency is scanned at 1 kHz. The resonances are jittering around due to mechanical noise in the piezo-actuated flexure stage holding the mirror. **(b)**: Fourier transform of the resonance position, without damping (top) and with damping (bottom, magnified 10 ×). Damping is provided by small pieces of sorbothane, and the amplitude of the oscillation was reduced by a factor ~ 50. The data behind this plot consist of 80 waveforms like the one shown on the left, each containing 200 resonances. The resonance positions of 10 consecutive waveforms were concatenated to increase the spectral resolution, and the 8 redundant spectra were then averaged. Inset: Mechanical spectrum of one of our Thorlabs NanoMax stages. Red dots: experimentally measured relative amplitude of the driven oscillations. Blue line: fit to a simple harmonic oscillator model (Eq. 3.36).

It exhibits a single resonance with the same frequency and with a similar Q-factor as in the driven oscillator case (inset). Thus we are confident that the observed jitter originates from the NanoMax stage's mechanical resonance, although the driving force could not be identified.

The detrimental impact of this resonance can be mitigated by the addition of some damping, in the form of sorbothane wedges inserted between the fixed and moveable parts of the stage. As more and more wedges were added, we repeated the previous experiments, and we observed that the amplitude of the resonance in the Fourier transform spectrum decreased. We were able to reduce the amplitude of the vibration by a factor close to 50, as shown in the final spectrum in Figure 3.12(b)(bottom). This proved to be enough for our needs, but it is not satisfactory in the long term. In particular it makes locking the cavity impossible. Eventually we will have to replace the flexure stage by a more compact and stable mount, currently under development, that can fit in a vacuum chamber or in a cryostat.

These studies on single microcavities were an important step in learning empirically how to manipulate, align and fine-tune both the waveguides and the spherical mirrors. The successful coupled cavities experiments described in the next chapter would not have been possible without this preliminary work.

Chapter 4

Coupled cavities

Armed with the basic understanding of lossy optical resonators developed in the previous chapter, we can now undertake to couple cavities together. The coupled-cavities system we envision, as described in Chapter 1.5.3, is a chain of many interconnected cavities. We start by considering the elementary unit of the chain: a free space microcavity coupled to a waveguide cavity. Later in this chapter we consider two coupled waveguide resonators with phase shifters. These provide the adjustable coupling between links of the chain.

4.1 Microcavity coupled to a waveguide cavity

The waveguide cavity-microcavity unit is depicted in Figure 4.1, and consists of a waveguide cavity of length L_{w} between mirrors of reflectivity R_{w} and R_{cw} and a free space microcavity of length L_{c} between mirrors R_{cw} and R_{c}. As before, R denotes the intensity reflectivity, and r the (real) amplitude reflection coefficient, $R = r^2$. We also define ϕ_{c} and ϕ_{w} as the microcavity and waveguide round trip phases. This configuration is reminiscent of a recent proposal (137), where two cavities are connected by a shared moveable mirror. By moving the mirror, a photon can be adiabatically transferred from one cavity to the other. The analysis developed in this paper, however, only applies to cavities of similar lengths whereas our waveguide resonator is about a hundred times longer than the microcavity.

We therefore follow a different approach. We start by deriving theoretical expressions for the complete reflection spectrum and for the coupling rate g_{cw}. We then demonstrate such coupling experimentally.

4.1.1 Derivation of the coupled-cavities spectrum

Lossless case

The most straightforward derivation of the coupled cavity spectrum is to consider the microcavity as a frequency-dependent effective mirror, with a complex amplitude reflection

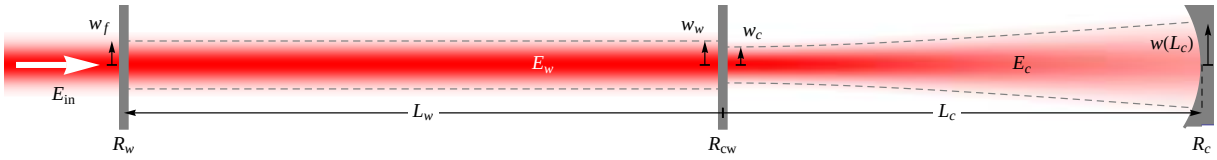

Figure 4.1: Schematics of the waveguide-microcavity system. The coupling fibre, waveguide and microcavity waists are w_f, w_w and w_c, and the mode-matching efficiencies are η_w^2 and η_cw^2. Waveguide and microcavity lengths L_c and L_w, mirror reflectivities R_w, R_cw and R_c

coefficient $\tilde{r}_\mathrm{c}e^{i\theta}$,and to substitute it for the output mirror in a simple cavity. The effective mirror reflectivity is given by Eq. (3.4):

$$\tilde{R}_\mathrm{c}(\omega) = \tilde{r}_\mathrm{c}^2 = \frac{r_\mathrm{cw}^2 + r_\mathrm{c}^2 - 2r_\mathrm{cw}r_\mathrm{c}\cos\phi_\mathrm{c}}{1 + r_\mathrm{cw}^2 r_\mathrm{c}^2 - 2r_\mathrm{cw}r_\mathrm{c}\cos\phi_\mathrm{c}} \tag{4.1}$$

$$\theta(\omega) = \arctan\left[\frac{(r_\mathrm{cw}^2 - 1)r_\mathrm{c}\sin(2\phi_\mathrm{c})}{r_\mathrm{cw}(r_c^2 + 1) - r_c(r_\mathrm{cw}^2 + 1)\cos(2\phi_\mathrm{c})}\right] \tag{4.2}$$

These can then be plugged back into Eq. (3.4) again, so that the field reflected from the coupled cavity is

$$E_r = E_{in} \times \frac{r_\mathrm{w} - \tilde{r}_\mathrm{c}e^{i(\theta+\phi_\mathrm{w})}}{1 - r_\mathrm{w}\tilde{r}_\mathrm{c}e^{i(\theta+\phi_\mathrm{w})}} \tag{4.3}$$

This produces a rich spectrum, a typical example of which is plotted in Figure 4.2(a). There are resonances whenever

$$\phi_\mathrm{w} + \theta = 0[2\pi]. \tag{4.4}$$

To think about this spectrum, let us start with both cavities independently on resonance ($\phi_\mathrm{w} = 2\pi m, \phi_\mathrm{c} = 2\pi p$) for an arbitrary reference light frequency $\Delta\omega = 0$ (which typically would also be the atomic transition).

- At the microcavity resonance ($\Delta\omega = 0$), the effective mirror introduces a phase shift $\theta = \pi$. Consequently the effective waveguide round trip phase is $\phi_\mathrm{w} = 2\pi(m + 1/2)$, i.e. it is exactly off-resonance. We see in Fig. 4.2(a) that there is indeed a reflected intensity dip at $\Delta\omega = 0$. On the other hand, the resonance condition is now satisfied at detunings $\Delta\omega = \pm g_\mathrm{cw}$ such that $\phi_\mathrm{w}(\pm g_\mathrm{cw}) + \theta(\pm g_\mathrm{cw}) = 2\pi m$, which gives rise to the splitting observed near zero detuning in Fig. 4.2(a). g_cw is the normal mode splitting, i.e. the photon tunnelling rate between the two cavities, as will be justified shortly.

- Away from the microcavity resonance, $\theta \simeq 0$ and the resonances are close to those of the waveguide alone. However, their visibility will be determined by the mismatch between the reflectivities of the input mirror R_w and of the effective mirror $\tilde{R}_\mathrm{c}(\omega)$. This gives rise to the observed visibility modulation. At some point we will have $R_\mathrm{w} = \tilde{R}_\mathrm{c}(\omega)$ and the visibility will reach unity, provided that the following condition

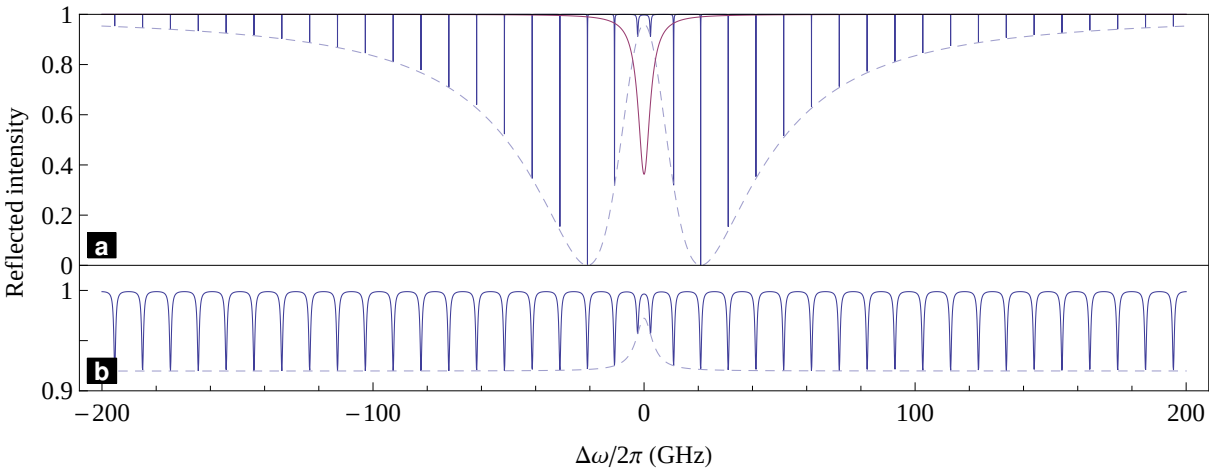

Figure 4.2: Spectrum of a waveguide-microcavity system (Eq. (4.3), solid blue curves) versus the detuning $\Delta\omega$ from the microcavity resonance. The waveguide and cavity lengths are set to 18700 and 100 wavelengths, respectively (10 mm and 78 μm), and the reflectivities are $R_{\rm w} = 99\%, R_{\rm cw} = 98\%, R_{\rm c} = 99.5\%$. **(a)** Ignoring waveguide propagation loss. **(b)** With 0.9 dB/cm loss. The pale, dotted curve is the envelope (Eq. (4.6)). The purple curve is the effective reflectivity $\tilde{r}_{\rm c}^2$ of the microcavity. The envelope reaches the x-axis when $\tilde{r}_{\rm c}^2(\Delta\omega) = R_{\rm w}$. The large phase shift introduced by this effective mirror causes the waveguide resonance at $\Delta\omega = 0$ to split: this is the signature of strong coupling. This pattern repeats itself with a periodicity given by the micro-cavity free spectral range (about 2 THz).

is fulfilled:

$$r_{\rm w}^2 > \left(\frac{r_{\rm cw} - r_{\rm c}}{1 - r_{\rm cw} r_{\rm c}}\right)^2 \tag{4.5}$$

The fringe envelope (dashed line) can be easily extracted by setting $e^{i(\theta+\phi_{\rm w})} = 1$ and is given by

$$I_{\rm min} = \frac{r_{\rm w} - \tilde{r}_{\rm c}}{1 - r_{\rm w}\tilde{r}_{\rm w}} \tag{4.6}$$

Tuning the microcavity length $L_{\rm c}$ essentially moves the envelop around, while tuning the waveguide length $L_{\rm w}$ moves the fringes under the envelopes. Eq. (4.4) cannot generally be solved analytically, but we can determine when the reflected intensity becomes zero. This is of interest because the spectrum is generally sensitive to the presence of atoms under this condition, as the field intensity in the microcavity is maximum. Zero reflected field is possible only if Eq. (4.5) is fulfilled, and it is achieved when the cavity can be tuned to make $r_{\rm w}^2 = \tilde{r}_{\rm c}^2(\phi_{\rm c})$. This occurs when

$$\phi_{\rm c}({\rm opt}) = \arccos\left[\frac{r_{\rm w}^2(1 + (r_{\rm c} r_{\rm cw})^2) - (r_{\rm c}^2 + r_{\rm cw}^2)}{2 r_{\rm cw} r_{\rm c}(r_{\rm w}^2 - 1)}\right]. \tag{4.7}$$

θ is then determined from Eq. (4.1) and $\phi_{\rm w}$ from (4.4).

Because the microcavity is much shorter than the waveguide, the effective mirror parameters $\tilde{r}_{\rm c}$ and θ can be considered constant around the resonances, whose width (ie

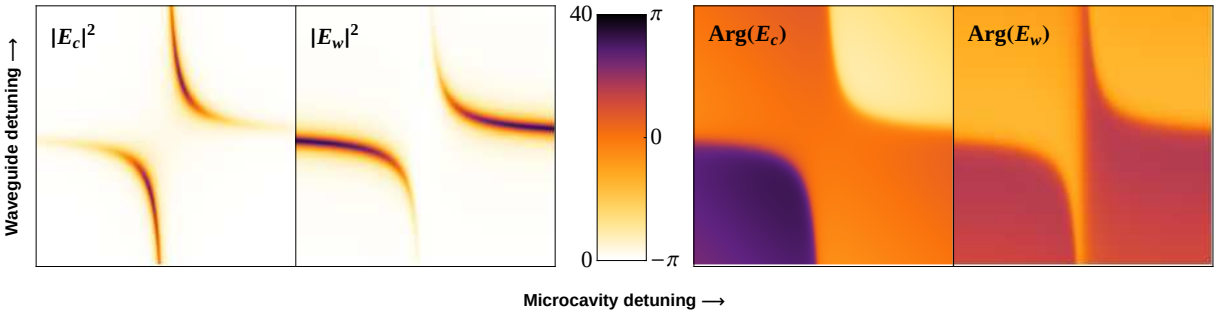

Figure 4.3: Intensity and phase of the fields circulating in the waveguide and in the microcavity, around the split resonance (equations (4.9)-(4.10)); reflectivities have been lowered for clarity compared to Fig. 4.2: $R = 0.90, R_{\mathrm{cw}} = 0.92, R_{\mathrm{c}} = 0.97$). The avoided crossing and the symmetric/antisymmetric nature of the two central resonances is evident, showing that the splitting is indeed the signature coherent energy exchange (or strong coupling) between the two cavities.

the coupled-cavity field decay rate) can therefore be approximated according to Eq. (3.11) as

$$\tilde{\kappa} = \frac{c(1 - r_{\mathrm{w}}\tilde{r}_{\mathrm{c}})}{2L_{\mathrm{w}}}. \tag{4.8}$$

The circulating field amplitudes inside the two cavities, E_{w} and E_{c}, can be easily computed in the effective mirror formalism. First we calculate the circulating amplitude in the waveguide E_{w} as in Eq. (3.13), using the effective mirror model of the microcavity. Then we use E_{w} as the input field for the microcavity. We obtain

$$E_{\mathrm{w}} = E_{\mathrm{in}} \frac{t_{\mathrm{w}}}{1 - r_{\mathrm{w}}\tilde{r}_{\mathrm{c}} e^{i(2\phi_{\mathrm{w}} + \theta)}} \tag{4.9}$$

$$E_{\mathrm{c}} = E_{\mathrm{w}} \frac{t_{\mathrm{cw}}}{1 - r_{\mathrm{cw}} r_{\mathrm{c}} e^{2i\phi_{\mathrm{c}}}} = E_{\mathrm{in}} \frac{t_{\mathrm{cw}} t_{\mathrm{w}} e^{i\phi_{\mathrm{w}}}}{1 - r_{\mathrm{c}} r_{\mathrm{cw}} e^{2i\phi_{\mathrm{c}}} + r_{\mathrm{c}} r_{\mathrm{w}} e^{2i(\phi_{\mathrm{c}} + \phi_{\mathrm{w}})} - r_{\mathrm{cw}} r_{\mathrm{w}} e^{2i\phi_{\mathrm{w}}}} \tag{4.10}$$

Figure 4.3 shows a map of the circulating field intensity and phase in the two resonators, $|E_{\mathrm{c}}|^2$, $|E_{\mathrm{w}}|^2$, $\arg(E_{\mathrm{w}})$ and $\arg(E_{\mathrm{c}})$, as a function of waveguide and microcavity detuning. The upper and lower branches correspond to the two peaks of the split resonance we previously described in Fig. 4.2. The laser frequency is set to $\Delta\omega = 0$. Here it is clear that the split resonance is actually an avoided crossing. The two branches are identical, except for the microcavity phases, which are opposite in sign. Thus we interpret the two resonances as the symmetric and antisymmetric normal modes of a coupled oscillator. It follows that, if the two normal modes are initially excited with the same amplitude, for example by placing a photon in the microcavity alone, the time evolution of the system will be a periodic exchange of the excitation between the two cavities, with a frequency g_{cw} corresponding to half the normal mode splitting.

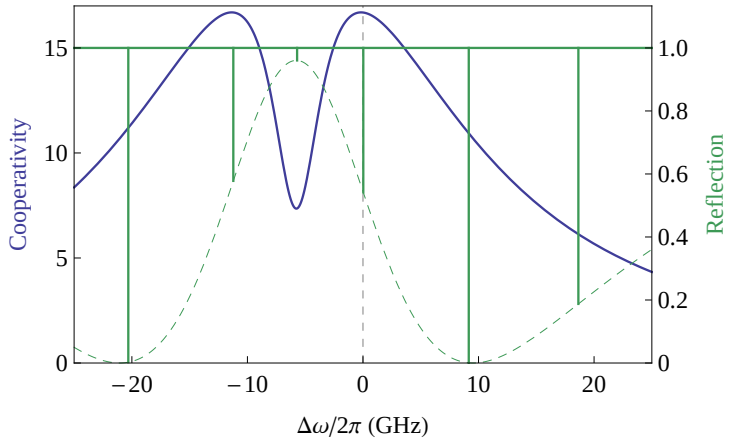

Figure 4.4: Cooperativity of a Rb atom in a microcavity coupled to a waveguide cavity ($R_\mathrm{w} = 0.98, L_\mathrm{w} = 23\,\mathrm{mm}, R_\mathrm{cw} = 0.98, L_\mathrm{c} = 156\,\mu\mathrm{m}, R_\mathrm{c} = 0.999$). Both cavities have been tuned to bring a resonance (green lines, dashed envelope) and maximise the cooperativity at the atomic frequency at $\Delta\omega = 0$.

Cooperativity

The cooperativity for an atom in the microcavity coupled to a waveguide resonator will be very similar to the single microcavity case of Eq. (3.35). The field decay rate κ is replaced by the coupled decay rate $\tilde{\kappa}$, and the atom-photon coupling g_ac of Eq. (1.7) is modified since the photon is now distributed between the two cavities according to the ratio $E_\mathrm{w}/E_\mathrm{c}$.

If the system is populated with a single photon of energy $\hbar\omega$, the field amplitudes become

$$\mathcal{E}_\mathrm{c} = \sqrt{\frac{\hbar\omega}{2\varepsilon_0\pi w_0^2}}\sqrt{\frac{|E_\mathrm{c}|^2}{L_\mathrm{c}|E_\mathrm{c}|^2 + L_\mathrm{w}|E_\mathrm{w}|^2}} \tag{4.11}$$

$$\mathcal{E}_\mathrm{w} = \sqrt{\frac{\hbar\omega}{2\varepsilon_0\pi w_0^2}}\sqrt{\frac{|E_\mathrm{w}|^2}{L_\mathrm{c}|E_\mathrm{c}|^2 + L_\mathrm{w}|E_\mathrm{w}|^2}} \tag{4.12}$$

We have assumed here that the microcavity mode is a Gaussian with constant waist, and that the waveguide and microcavity waists are identical: $w_\mathrm{c} = w_\mathrm{w} = w_0$. A full expanding Gaussian treatment similar to the one presented in section 3.5.1 is possible, but would result only in minor quantitative differences. The atom-photon coupling is then given by the fraction of the total field that resides in the microcavity:

$$g'_\mathrm{ac} = g_\mathrm{ac}\sqrt{\frac{|E_\mathrm{c}|^2}{|E_\mathrm{c}|^2 + \frac{L_\mathrm{w}}{L_\mathrm{c}}|E_\mathrm{w}|^2}}. \tag{4.13}$$

The cooperativity is then $C = {g'_\mathrm{ac}}^2/2\tilde{\kappa}\gamma$, and is plotted in Figure 4.4. The double-bump structure is caused by the interplay between g'_ac and $\tilde{\kappa}$, both of which peak at $\phi_\mathrm{c} = 0$ (microcavity on resonance) whereas $C \propto g^2/\kappa$. By detuning both the microcavity

(which translates the whole cooperativity curve, very much like the visibility function) and the waveguide cavity, it is possible to bring a resonance at the desired frequency (here $\Delta\omega = 0$) while maximising the cooperativity, as we have done in Fig. 4.4.

It should be noted that maximum cooperativity does not coincide with maximum visibility, as we had anticipated (see Eq. (4.7)). This is because at this point the waveguide field is also rather large and the photon wavepacket is not strongly localised in the microcavity.

Lossy waveguide

We can include the waveguide propagation loss, mirror loss and mode mismatch exactly as we did in the previous chapter (see section 3.2). Figure 4.2(b) shows the same coupled cavity as before, but with a waveguide propagation loss of 0.9 dB/cm. We observe a similar splitting, since the physics of the system is identical, but the spectrum looks less dramatic. Besides the inevitable broadening of the resonances, the main difference is that the fringe visibility never reaches unity, because Eq. (4.5) (with r_{cw} effectively reduced by the loss) can no longer be fulfilled.

Mode mismatch and mirror loss do not qualitatively alter this spectrum. Their main effect, in the relevant operating conditions, is to significantly decrease the visibility and increase the splitting g_{cw}. They need to be taken into account when modelling a realistic system. We do not reproduce here the complicated expressions for the reflectivity; they can be easily derived from the previous discussion.

4.1.2 Analytical expression of the coupling rate g_{cw}

We would very much like to derive an analytical expression for the coupling rate g_{cw} between the two cavities. However it is not possible to do so directly from Eq. (4.4) because the arctan in Eq. (4.2) cannot be linearised over the required range. Fortunately there is a different approach to the problem, based on rate equations and developed by Lang and Yariv in (138). It does not allow us to derive the whole coupled-cavity spectrum as easily as we have done, but it does provide relatively simple analytical expressions for the coupling. Through a somewhat cumbersome procedure, we obtain a set of coupled rate equations for the freely evolving field amplitude in each cavity, E_{c} and E_{w}:

$$dE_{\mathrm{c}}(t)/dt = i\kappa_{\mathrm{cc}}E_{\mathrm{c}}(t) + i\kappa_{\mathrm{wc}}E_{\mathrm{w}}(t) \tag{4.14}$$

$$dE_{\mathrm{w}}(t)/dt = i\kappa_{\mathrm{ww}}E_{\mathrm{w}}(t) + i\kappa_{\mathrm{cw}}E_{\mathrm{c}}(t) \tag{4.15}$$

with the complex cross- and self-coupling coefficients

$$\begin{aligned} \kappa_{\mathrm{ww}} &= \tfrac{c}{2iL_{\mathrm{w}}}(r_{\mathrm{w}}r_{\mathrm{cw}}e^{i\phi_{\mathrm{w}}} - 1) & \kappa_{\mathrm{cc}} &= \tfrac{c}{2iL_{\mathrm{c}}}(r_{\mathrm{c}}r_{\mathrm{cw}}e^{i\phi_{\mathrm{c}}} - 1) \\ \kappa_{\mathrm{wc}} &= \tfrac{-c}{2iL_{\mathrm{w}}}\eta_{\mathrm{cw}}t_{\mathrm{cw}}r_{\mathrm{w}}e^{i\phi_{\mathrm{w}}} & \kappa_{\mathrm{cw}} &= \tfrac{-c}{2iL_{\mathrm{c}}}\eta_{\mathrm{cw}}t_{\mathrm{cw}}r_{\mathrm{c}}e^{i\phi_{\mathrm{c}}} \end{aligned} \tag{4.16}$$

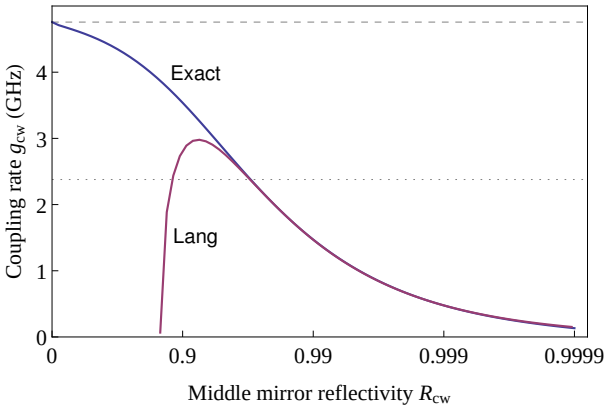

Figure 4.5: Waveguide-microcavity coupling as a function of middle mirror reflectivity R_{cw}. Blue curve: $g_{\mathrm{cw}}^{\Delta\omega}$, the spectral splitting calculated numerically from Eq. (4.4. Purple curve: g_{cw}^{L}, Lang's analytical solution of Eq. (4.19). Other parameters are $R_{\mathrm{c}} = 0.999$, $R_{\mathrm{w}} = 0.99$, $L_{\mathrm{c}} = 156\,\mu\mathrm{m}$, $L_{\mathrm{w}} = 100L_{\mathrm{c}}$. The dashed line indicates the FSR of a cavity of length $L_{\mathrm{c}} + L_{\mathrm{w}}$, and the dotted line to half the FSR of a cavity of length L_{w}.

t_i, r_i are the usual amplitude transmission/reflection coefficient and η_{cw} is the mode-matching efficiency.

The solutions of Eq. (4.14) are damped sinusoidal oscillations of the form

$$E_{\mathrm{c}}(t) \quad \propto e^{it(\kappa_{\mathrm{cc}}+\kappa_{\mathrm{ww}})/2}\sin(tg_{\mathrm{cw}}) \tag{4.17}$$

$$g_{\mathrm{cw}} \quad = \tfrac{1}{2}\sqrt{(\kappa_{\mathrm{cc}}+\kappa_{\mathrm{ww}})^2 + 4\kappa_{\mathrm{cw}}\kappa_{\mathrm{wc}}} \tag{4.18}$$

Assuming that the outside mirrors have a high reflectivity ($r_{\mathrm{c}} \simeq r_{\mathrm{w}} \simeq 1$), the coupling rate can be rewritten as

$$g_{\mathrm{cw}} = \frac{c}{4L_{\mathrm{c}}M}\sqrt{M(1-r_{\mathrm{cw}})}\sqrt{8 - M(1-r_{\mathrm{cw}}) + 8}. \tag{4.19}$$

$M \gg 1$ is the ratio of the cavity lengths $L_{\mathrm{w}}/L_{\mathrm{c}}$. g_{cw} of Eq. (4.18) (hereafter denoted g_{cw}^{L}) is plotted in Figure 4.5 (purple curve), together with the spectral splitting calculated by numerically solving the resonance condition of Eq. (4.4) (blue curve, hereafter $g_{\mathrm{cw}}^{\Delta\omega}$).

For $R_{\mathrm{cw}} \to 1$, there can be no energy exchange between the cavities and $g_{\mathrm{cw}} \to 0$. Then as R_{cw} decreases, the splitting increases as photon tunnelling is turned on; (4.18) and (4.4) agree very well down to $R_{\mathrm{cw}} \simeq 0.93$, which is sufficient for our purpose (we envision $R_{\mathrm{cw}} > 0.98$ at minimum).

At very low reflectivity, it is as if there was no middle mirror at all: the spectral splitting $g_{\mathrm{cw}}^{\Delta\omega}$ becomes the half-FSR of a cavity of length $L_{\mathrm{c}} + L_{\mathrm{w}}$ (dashed horizontal line). Lang's solution on the other hand becomes imaginary for $R_{\mathrm{cw}} < (1 - 8/M)^2$, which roughly coincides with g_{cw} becoming larger than $\mathrm{FSR}_{\mathrm{w}}/4$ (dotted line). This is because the coupled equations (4.14) assume adiabaticity, and this assumption breaks down when the coupling time scale becomes comparable to the photon round trip time of the longer cavity.

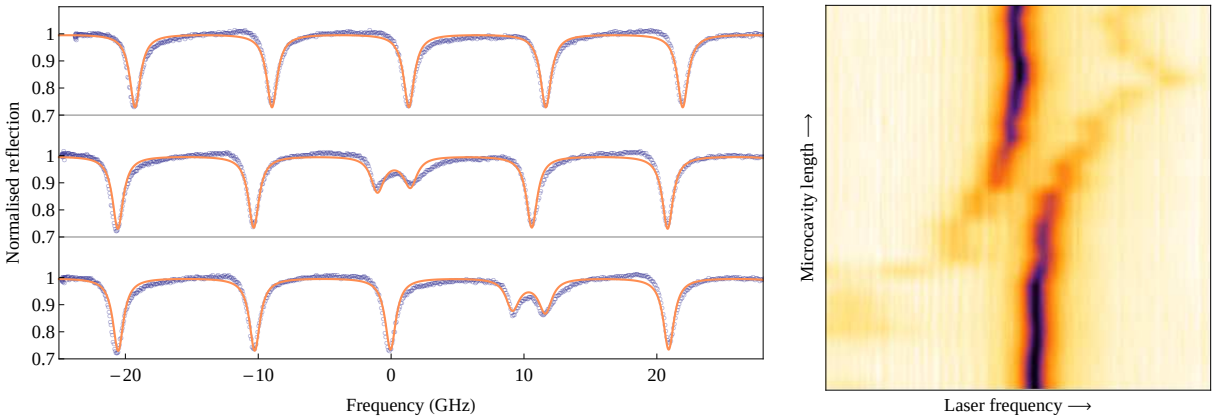

Figure 4.6: Observation of waveguide-microcavity coupling on ORC7. Top: the microcavity is far detuned and we only observe the waveguide spectrum. Middle and bottom: the microcavity is tuned to two different waveguide resonances. The solid lines are fits to Eq. (4.3), including waveguide and mode matching losses. Right: scanning the laser frequency and the microcavity length reveals that the observed splitting is a proper avoided crossing. The frequency and length axes span about 8 GHz and 50 nm respectively (the exact calibration is not available).

4.1.3 Experiments

Initial waveguide-microcavity coupling experiments were carried out on chip "CIP", but the loss proved too large, and although the effect of the microcavity on the waveguide cavity could be observed as a small visibility reduction, we could not access the strong coupling regime. We then moved on to the "ORC7" chip with mirror reflectivities $R_{\mathrm{w}} = 90$ and $R_{\mathrm{cw}} = 99\%$. Results are presented in Figure 4.6. Top left we see the waveguide cavity spectrum alone, the microcavity being far detuned. We fit the full cavity model of Eq. (3.4) including propagation loss (0.9 dB/cm) and fibre-waveguide mode mismatch, from which we estimate the mode matching efficiency as $\eta_{\mathrm{w}}^2 = 40 \pm 2\%$. This is much lower than what we expected from the mode profile measurement ($\sim 90\%$, see Chapter 2).

Then we bring the microcavity into resonance by changing its length, and we observe the middle spectrum, exhibiting a clear splitting. We can tune the microcavity to the next waveguide resonance as well (bottom). Again we use a fit based on Eq. (4.3), including the parameters determined in the off-resonant measurement, to estimate the waveguide-microcavity mode-matching efficiency ($\eta_{\mathrm{cw}}^2 = 52 \pm 2\%$) and the spherical mirror reflectivity ($R_{\mathrm{c}} = 99.8 \pm 0.1\%$). R_{c} was not known precisely beforehand, but was expected to be around 99.9%. The observed splitting is reproduced fairly well by the analytical formula (4.19), which gives $2g_{\mathrm{cw}}^L/2\pi = 2.6$ GHz with the parameters determined by the fit. The splitting measured directly on the middle and bottom spectra is 2.7 ± 0.2 and 2.3 ± 0.2 GHz respectively. Here too the mode-matching efficiency is much lower than anticipated (see Fig. 3.9).

These disappointing mode-matching efficiencies could simply be an artifact of having neglected the mirror loss A_1. We have derived in Chapter 3.2 expressions for the visibility

in the presence of mirror or mode mismatch losses (Equations (3.16, 3.21)). It is relatively easy to combine both: the $(1-l)$ or $(1-A)$ terms are replaced by $1-[t_1^2(1-\eta^2)+A\eta^2]$. Thus introducing a mirror loss will effectively increase the mode-matching efficiency; however there is no way of distinguishing between the two. Plugging in some numbers, we see that we introduce a middle mirror loss $A_{\mathrm{cw}} \simeq 0.002$, η_{cw}^2 is increased to around 80%, which is sensible. However achieving a similar result for η_{w} requires $A_{\mathrm{w}} \simeq 2\%$: this is simply too high. Thus mirror loss can only be a part of the puzzle. We did not have the time to carry out similar experiments on the latest chip, ORC8.

4.2 Coupled waveguide cavities

Having demonstrated the coupling between waveguide and microcavity, we now move on to the second basic unit: two coupled waveguides. From a theoretical point of view the situation is somewhat reversed, in that it will be very easy to derive an analytical expression for the coupling rate g_{ww} and rather more difficult to work out the full spectrum. A similar although slightly more simple device was initially proposed by Yariv (139) in the context of coupled-resonator optical waveguides (CROWs) (140). Essentially, a coupled waveguide resonator array can be thought of as converting a space-like coupling (without the mirrors, propagating energy is exchanged between and along the waveguides) to a time-like coupling (with the mirrors, it is the circulating energy that is exchanged).

4.2.1 Analytical derivation of the coupling rate g_{ww}

We have seen previously (Chapter 2.4.1) that an array of N coupled waveguides is described in terms of N normal modes, each having a distinct k-vector β_i given by Eq. (2.10). Figure 4.21(a) shows the electric field amplitude profiles of the first eight normal modes of a coupled array of nine waveguides; the corresponding linear dispersion relations have been plotted in Fig. 4.21(b) (k-vectors vs optical frequency ω). What then happens if the waveguides are terminated by two mirrors to form cavities is illustrated by 4.7(b) and (c). It is essentially identical to the single cavity case: each normal mode has to satisfy the resonance condition

$$\beta_i L = m\pi. \tag{4.20}$$

Thus β_i is fixed and common to all normal modes (black vertical line in Fig. 4.21(b)), so that the resonance condition will be satisfied at a different frequency ω_i for different modes, for a given order m. Note that the waveguide length L can be larger than the coupler length l. For a homogeneous waveguide array (all waveguide-waveguide coupling constants κ_i and all waveguide optical lengths identical), applying the resonance condition

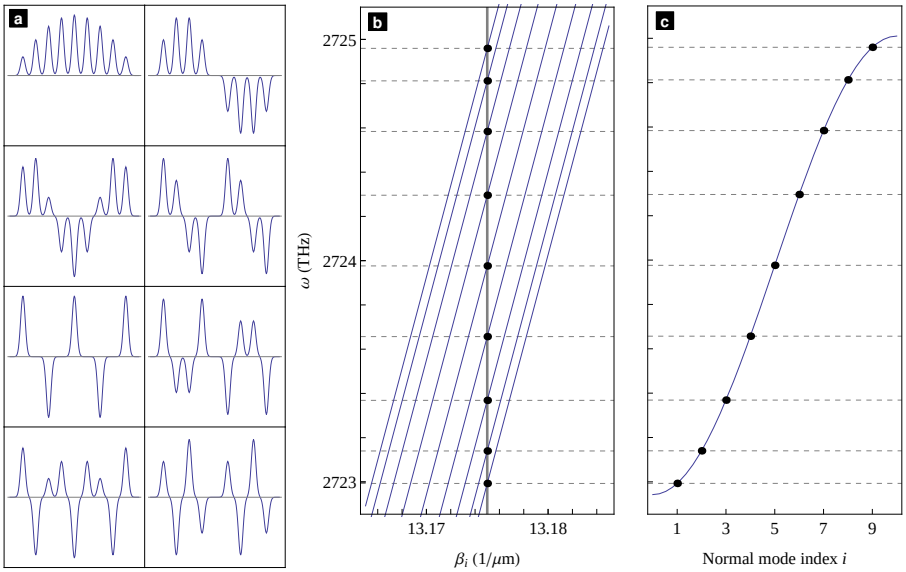

Figure 4.7: Illustration of the conversion from space-like to time-like coupling in a waveguide cavities array. **(a)** Electric field of the normal modes in a 9-waveguides array. **(b)** Each mode has a distinct propagation constant β_i, which is proportional to the frequency ω. The quantisation of β_i by a cavity (vertical line) results in different resonance frequencies ω_i for each normal mode **(c)**. [Adapted from (139)].

to Eq. 2.10 gives the N resonant frequencies

$$\omega_i = \omega_0 \left(1 - 2\frac{\kappa_l l}{m\pi} \cos \frac{i\pi}{N+1}\right) \tag{4.21}$$

as plotted in Figure 4.7(c). ω_0 is the resonant frequency of a single, uncoupled waveguide resonator.

The special case $N = 2$ is most interesting because the full splitting $\Delta\omega$ between the two normal modes gives immediately the photon tunnelling rate g_{ww} between the waveguides:

$$g_{\text{ww}} = \frac{\Delta\omega}{2} = \frac{\kappa_l l \omega}{m\pi} = \frac{\kappa_l l}{L} \frac{c}{n_{\text{eff}}} \tag{4.22}$$

We notice immediately that this expression only involves the product $\kappa_l l$, which can be easily measured prior to coating the mirrors since it related to the the splitting ratio $C = \sin^2(\kappa_l l)$. Thus we do not require any great knowledge of the inner workings of the coupler, which is fortunate as both κ_l and l are difficult to access on their own.

4.2.2 Derivation of the coupled cavity spectrum

The most intuitive way of describing a system of two coupled waveguide resonators follows the approach used to derive the single-cavity spectrum in Section 3.2. Refer to Figure 4.8

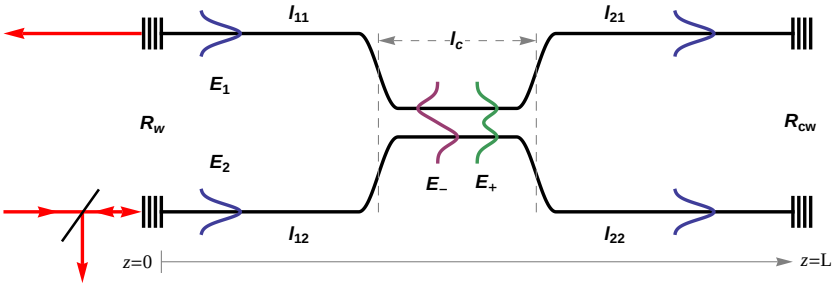

Figure 4.8: Two coupled waveguides cavities. In the four arms (whose lengths l_{ij} are not necessarily equal), a single mode propagate with k-vector $k = 2\pi n_{\text{eff}}/\lambda$. In the coupler section (length l_c), the symmetric and antisymmetric normal modes E_+ and E_- propagate with k-vector k_+ and k_-

.

for the notation. We describe the field amplitude by a vector

$$\mathbf{E}(z) = \begin{pmatrix} E_{1,+}(z) \\ E_{2,-}(z) \end{pmatrix}. \tag{4.23}$$

where E_1 and E_2 are the field amplitudes in the upper and lower waveguides, in the uncoupled section, while E_+ and E_- are the field amplitudes of the symmetric and antisymmetric normal modes in the coupler section. We switch between the $(+,-)$ and $(1,2)$ bases by multiplying $\mathbf{E}$ by the matrix

$$\mathbf{B} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \tag{4.24}$$

Propagation from the input mirror (at $z = 0$) to the coupler, and from the coupler to the output mirror at $z = L$, is represented by the two matrices

$$\mathbf{M}_{\text{in}} = \begin{pmatrix} e^{ikl_{11}} & 0 \\ 0 & e^{ikl_{12}} \end{pmatrix} \quad \text{and} \quad \mathbf{M}_{\text{out}} = \begin{pmatrix} e^{ikl_{21}} & 0 \\ 0 & e^{ikl_{22}} \end{pmatrix} \tag{4.25}$$

while reflection is taken care of by two more matrices $\mathbf{r}_{\text{w}} = r_{\text{w}}\mathbb{I}$ and $\mathbf{r}_{\text{cw}} = r_{\text{cw}}\mathbb{I}$. $\mathbb{I}$ is the identity matrix. The coupler is represented as

$$\mathbf{M}_{\text{coupler}} = \mathbf{M} \cdot \begin{pmatrix} e^{ik_+ l_c} & 0 \\ 0 & e^{ik_- l_c} \end{pmatrix} \cdot \mathbf{M}. \tag{4.26}$$

In this way it is quite straightforward to derive a round trip matrix $\mathbf{M}_{\text{rt}} = \mathbf{M}_{\text{in}} \cdot \mathbf{B} \cdot \mathbf{M}_{\text{coupler}} \cdot \mathbf{B} \cdot \mathbf{M}_{\text{out}} \cdot \mathbf{r}_{\text{cw}} \cdot \mathbf{M}_{\text{out}} \cdot \mathbf{B} \cdot \mathbf{M}_{\text{coupler}} \cdot \mathbf{B} \cdot \mathbf{M}_{\text{in}} \cdot \mathbf{r}_{\text{w}}$ (the lengthy expression does not need to be reproduced here). The intracavity field, close to the input mirror, in the (1,2)

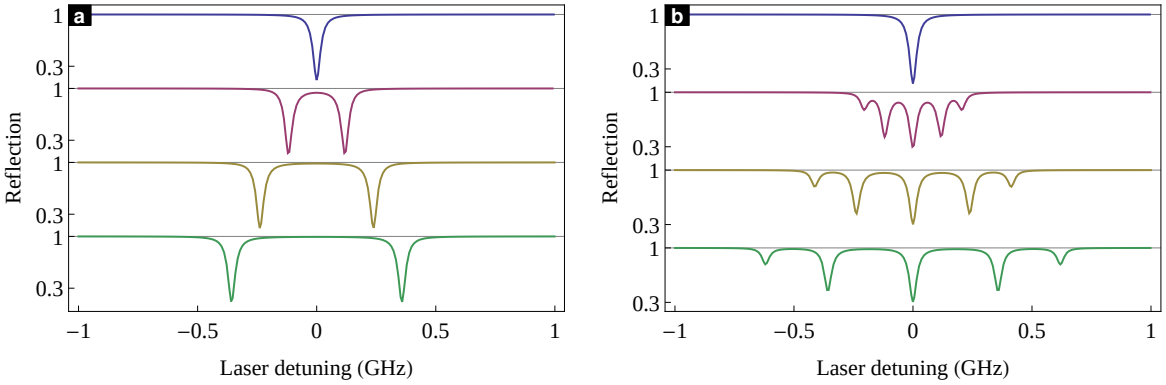

Figure 4.9: Reflection spectrum from the input waveguide of a two- and five-waveguide coupled resonator array, for various values of the coupling constant $\kappa_l = 0, 500, 1000, 1500\,\mathrm{m}^{-1}$. The mirror reflectivities are 98 and 99%.

basis, assuming we are pumping only one of the waveguides, is given by the infinite sum

$$\mathbf{E}_{\mathrm{intra}} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \cdot \sqrt{\mathbb{I} - \mathbf{r}_{\mathrm{w}}^2} \cdot \sum_{n=0}^{\infty} \mathbf{M}_{\mathrm{rt}}^n. \tag{4.27}$$

$\sqrt{\mathbb{I} - \mathbf{r}_{\mathrm{w}}^2}$ is the matrix field transmission coefficient through the input mirror. We can then use the following identity (convergence is trivial for physical reasons):

$$\sum_{n=0}^{\infty} \begin{pmatrix} a & b \\ c & d \end{pmatrix}^n = \frac{1}{a + bc + d - ad - 1} \begin{pmatrix} d-1 & -b \\ -c & a-1 \end{pmatrix} \tag{4.28}$$

and derive a closed form, if complicated, expression for the coupled-waveguide cavity spectrum. A typical example is shown in Figure 4.9(a). We see two resonances indeed, whose splitting is linear in κ_l and reproduces exactly Eq. (4.22).

In order to strengthen our interpretation of the splitting as a coupling rate, we plot in Figure 4.10 the intracavity intensities and phases as a function of both waveguide detunings, exactly as we did earlier for a waveguide coupled to a microcavity (Fig. 4.3). Again, we see a clear avoided crossing and the two branches (corresponding to the two resonances of Fig. 4.9(a)) are identical, except for the phase sign in the second waveguide. Thus the two branches can be interpreted as the symmetric and anti-symmetric normal modes, and the frequency splitting between them as a coupling rate describing the rate of energy exchange between the waveguides.

Note that this method can be easily extended to larger arrays, since the formula corresponding Eq. (4.28) can be extended to arbitrary large matrices. The spectrum of an array of five waveguides has thus been plotted in Figure 4.9(b). However, such analytical expressions quickly become extremely long and cumbersome. For larger problems, a numerical method based on transfer matrices can be used, such as that presented in (139) (with some minor adaptations). This requires a matrix inversion for every point to be

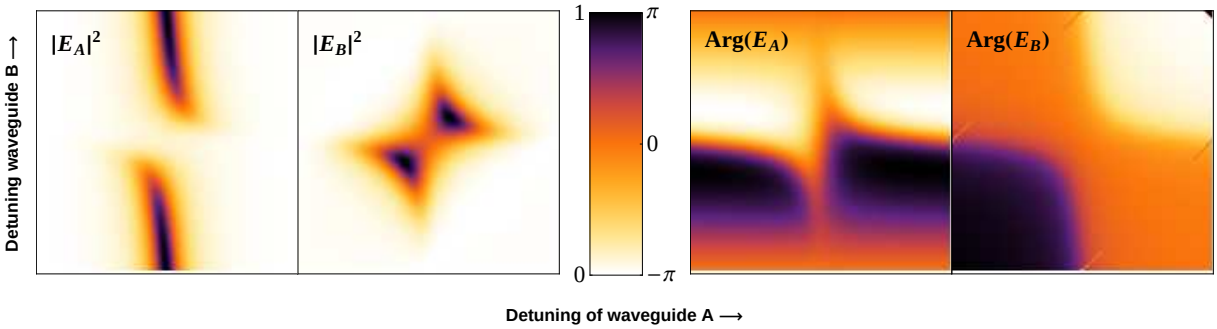

Figure 4.10: Intensity and phase of the fields circulating in two coupled waveguides when changing the lengths of (detuning) the two waveguides. The plots cover $\pm\lambda/4$ on both axis, ie a full period. A is the input waveguide. The plots are very similar to the waveguide-microcavity case (Fig. 4.3) and illustrate the symmetric/antisymmetric nature of the normal modes. The avoided crossing is clearly resolved.

calculated and therefore becomes slow for large arrays. Since we limit ourselves to arrays of two waveguides, we'll stick to the analytical calculation, which has been used to make the theoretical plots that follows in Figures 4.15 to 4.17.

4.2.3 Tuning the coupling rate with phase shifters

When describing the operation of evanescent couplers in Chapter 2.4.1, we showed that the coupling constant κ_l is proportional to the overlap integral between the two modes (Eq. (2.5)). In a coupled-resonator waveguide array, the overlap integral must also account for the standing-wave nature of the modes, as illustrated in Fig. 4.11. It is straightforward to show that this introduces a phase factor such that κ_l becomes $\kappa_l \cos\phi$. ϕ is just the phase difference between the two standing waves: $\cos\phi = 1$ (full coupling) if they are in phase (nodes facing each other), and $\cos\phi = 0$ (no coupling) if they are in quadrature.

Therefore, to use a simple image, all we need to do to tune the coupling is to translate one of the standing waves up and down the waveguide. This we can easily do by inserting phase shifters on both sides of the coupler, and tuning them so as to get equal magnitude but opposite sign phase shifts to keep the cavity on resonance. Using NiCr heaters as phase shifters, a negative phase shift means cooling, so that we have to operate the heaters around some hot setpoint.

4.2.4 Experiments

Jena chips higher-mode coupling

The first observation of waveguide-waveguide cavity coupling was entirely accidental, on a chip ("Jena") that was not intended for coupling but did contain closely spaced waveguides ($\sqrt{2}.25\,\mu$m $= 35\,\mu$m apart; see Figure 2.7). We had coupled the input fibre to one of the waveguide, but the output fibre had been coupled to the next waveguide in line,

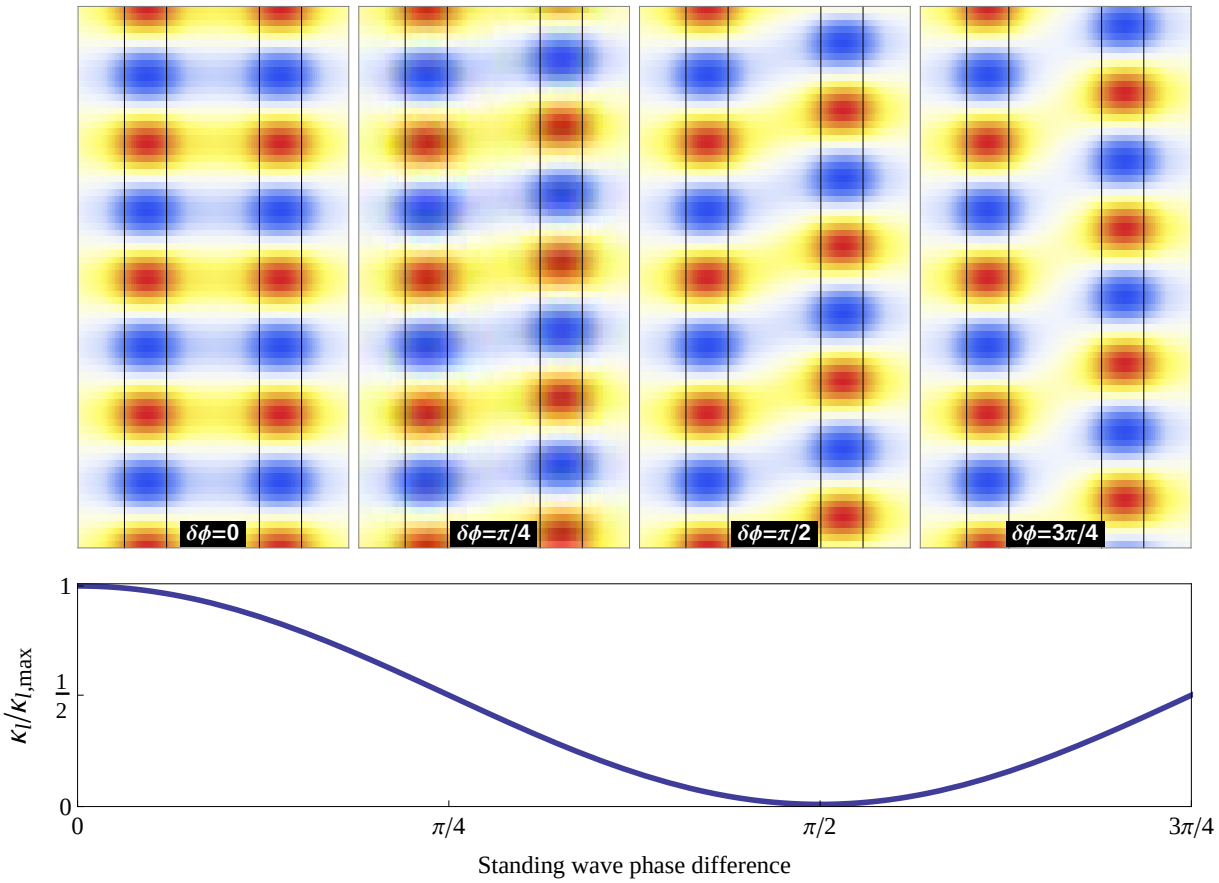

Figure 4.11: Modal amplitude in two evanescently coupled waveguides cavities (negative in blue, positive in red), and evanescent coupling constant $\kappa_l \cos^2(\delta\phi)$ as a function of phase difference $\delta\phi$ between the two standing waves.

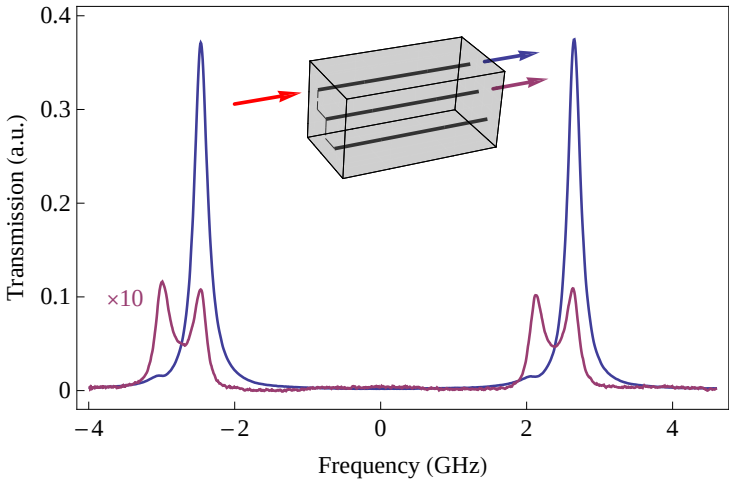

Figure 4.12: Transmission spectrum of the Jena femtosecond-written chip with end facet mirrors. Blue: direct transmission of the TEM_{00} mode. The finesse is 24. Purple: output of the second waveguide, 35 μm away from the first, which result from evanescent coupling of 0.2% of the TEM_{01} mode over a single pass.

and we observed that the transmission spectrum had an unexpected doublet structure (Figure 4.12, purple curve). The inset graphic clarifies the configuration. After realising the mistake, we moved the output fibre to the correct waveguide, and observed the blue spectrum, which looks more like what was expected from a single waveguide cavity. But what was going on in the second waveguide?

The first clue is the small bump we see in the main (blue) spectrum immediately to the left of the resonance. This seems to indicate the presence of a higher-order waveguide mode, whose amplitude is much smaller because its overlap with the coupling fibre mode is smaller. Moreover this bump coincides with the left peak of the second (red) spectrum: could they be one and the same feature? To confirm this we quickly looked at the mode profile in the second waveguide. It was shown to have a somewhat blurry double-lobe structure compatible with a TEM_{01} mode, with dimensions $7 \times 9\,\mu m^2$, significantly larger than the fundamental mode of the first waveguide ($5 \times 7\,\mu m^2$).

So the most likely explanation is this: because the higher-order mode is wider, it couples evanescently more strongly than the fundamental. Therefore, although both modes co-propagate in the first waveguide, only the higher-order one couples to the second waveguide and experiences the observed splitting. To test this hypothesis quantitatively, we use Eq. (4.22), which relates the splitting $\Delta\omega = 2\pi \times 0.54\,\text{GHz}$ and the coupling constant κ_l, which we evaluate from Eq. 2.5 based on the mode size given above. The only free variable left is the waveguide index contrast Δn, which we do not know. But we can adjust it so as to fulfil Eq. (4.22), which yields $\Delta n \sim 2\,10^{-3}$. This is a perfectly sensible value for femto-second written waveguides.

We can also calculate the single-pass coupling ratio, which is given by Eq. 4.22 as $\sin^2(l\kappa_l) = 0.2\%$. It is therefore not surprising that we never detected this coupling before

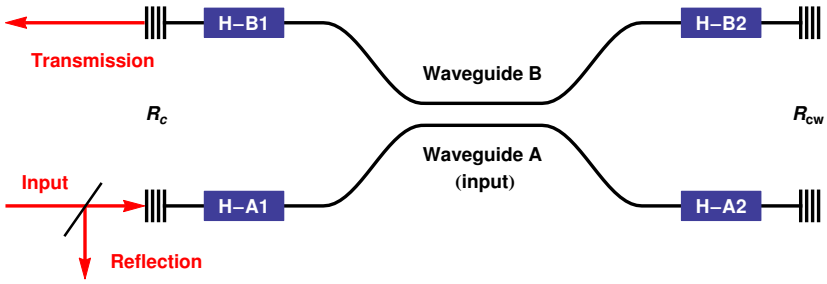

Figure 4.13: Schematics and notation for the coupled waveguide resonators experiments. The probe laser is injected into waveguide A, and we monitor the "reflection" from waveguide A and the "transmission" from waveguide B. H-A1 etc. are the four thermo-optic phase shifters.

depositing the mirrors.

ORC8

The final waveguide-waveguide coupling experiments were carried out on chip "ORC8", which contained five evanescent and five cross-couplers (as well as four straight waveguides). The splitting ratios had all been measured prior to fabricating the mirrors (see Fig. 2.13). After the mirror deposition, most of the couplers showed evidence of splitting. Since we could only fit six sets of heaters on the chip, we selected the most promising couplers (highest finesse, best symmetry). The finesse was in the range 7–16. We then deposited the NiCr heaters, and we were able to tune all six selected couplers into resonance, to observe anticrossing, and to tune the splitting. However quite a few of the heaters broke early on (because of fabrication defects or of human error), so that we were unable to record usable data on two couplers, and only a limited set on others.

The sketch in Figure 4.13 clarifies the notation used in the following. Laser light is butt-coupled from a v-groove fibre array into one arm (A) of a coupler. We monitor both the reflection from waveguide A, and the "transmission" from the second waveguide (B), which is collected in the next fibre of the v-groove array. The four heaters are denoted H-A1, H-A2, H-B1 and H-B2: "1" refers to heaters on the input side of the coupler, "2" to the opposite side.

In a first experiment we used a 800 μm long evanescent coupler ("EVA-2"), with a measured travelling-wave splitting ratio $C_\times = 18 \pm 2\%$. We scanned the laser frequency while increasing the power dissipated into one of the input heaters (H-A1), and we recorded the reflection spectrum from the same waveguide. In Figure 4.14, we can see the resonance fringes move to lower frequencies, very much like Fig. 3.6, until it gets close to the resonance of the second waveguide (up to this point a barely noticeable bump). This second resonance then gains amplitude at the expense of the first one, in what looks like an avoided crossing. The small shift of this second resonance is caused by thermal cross-talk.

To demonstrate anticrossing further, we set the two waveguides on resonances, switched

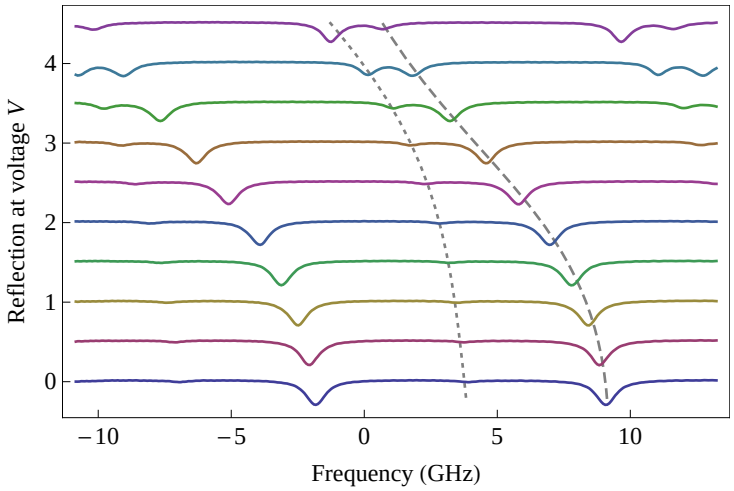

Figure 4.14: Evolution of the reflection spectrum as the the voltage in heater H-A1 is increased, showing avoided crossing between the resonances of the two waveguides (A: dashed, B: dotted). [ORC8-EVA2]

off the laser scan and adjusted the laser frequency to sit right in the middle of the split resonance. We then modulated the power dissipated in the input heater H-A1 and its opposite number on the second waveguide H-B1 simultaneously. H-A1 was modulated at 40 Hz and H-B1 at 0.5 Hz so that a full scan took only 2 s. Both heaters were calibrated separately to convert power into a frequency detuning. The results are plotted in Figure 4.15 for both the reflection and transmission, together with a theoretical calculation (inset).

When the two waveguides are non-resonant, scanning H-A1 and observing the reflection (A) only reveals a single resonance fringe, that of the input waveguide A. Then as H-B1 is tuned and cavity B becomes resonant with cavity A, we see a splitting which develops into an avoided crossing. This demonstrates strong coupling between the two resonators. The tilt of the dark band in (A) is caused by thermal crosstalk from H-B1.

Detuning the two heaters by the same amount ($\delta f_{\mathrm{H-B1}} = \delta f_{\mathrm{H-A1}}$) is equivalent to detuning the laser; thus by measuring the resonant splitting along the dashed grey line at 45° one can estimate the coupling rate as $g_{\mathrm{ww}} = 0.65 \pm 0.05\,\mathrm{GHz}$. This is quite significantly lower than the theoretical expectation given by Eq. 4.22, with the experimentally measured splitting ratio $C_\times = \sin^2(\kappa_l l) = 18 \pm 2\%$, which gives $g_{\mathrm{ww}} = 1.4 \pm 0.1\,\mathrm{GHz}$. But there is no reason to expect the maximum coupling, since we have yet to tune the standing waves relative phase.

This is what we are about to do now, on a cross coupler this time ("X-2", crossing angle 2.3°, splitting ratio $C_\times = 32 \pm 2\%$). Our control knobs are the two heaters on the second waveguide, H-B1 and H-B2, which were calibrated beforehand. We start by bringing waveguide cavity B into resonance with A. Then we manually increase the power in H-B1 in small steps, all the while balancing it by decreasing the power in H-B2 and recording frequency spectra. The results are plotted in Figure 4.16. It shows the

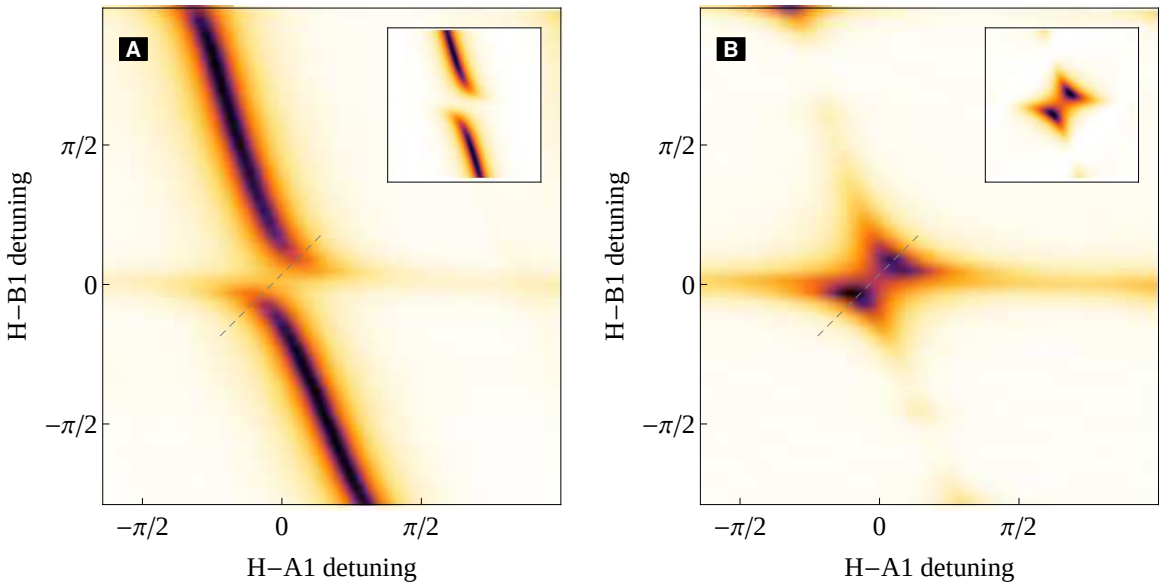

Figure 4.15: Avoided crossing on evanescent coupler ORC8-EVA2. It shows the reflection from waveguide A and transmission from waveguide B as the power in two phase shifters (one on each waveguide, specifically H-A1 and H-B1) is changed to detune the cavities. The avoided crossing is clearly resolved, demonstrating that the two resonators are strongly coupled. The data is very well reproduced by theoretical calculations (inset). The slow drift of the main resonance (whose frequency ought to be independent from H-A1) is caused by thermal cross talk from heater H-B1.

transmission out of waveguide B on the left and reflection from waveguide A on the right as a function of laser frequency and heater phase shift. Starting from the bottom, where there is no coupling, we can see the resonance broaden and split. The splitting reaches a maximum, and would reach zero coupling again after a phase shift of π. Theoretical calculations are shown in inset, and reproduce the data fairly well. The small overall drift is likely caused once again by thermal cross-talk. Note that the transmission signal does not vanish completely at $\phi = -\pi/2$, as we would expect since there should be no power in waveguide B. This is only because of the experimenter's inadequacy: manual tuning of the phase shifters becomes tricky when the splitting is not resolved,

The final data (Figure 4.17) addresses this point. It is very similar to the previous one, but comes from a different coupler (cross-coupler "X-1", crossing angle 2.2°, splitting ratio $C_\times = 15 \pm 2\%$). Here, instead of manually balancing the two heaters, we automated a full scan of the three-dimensional parameter space (H-B1, H-B2, frequency). Out of the large data set, we extracted the spectra corresponding to a balanced heater tuning (ie, those that would be lying on the grey dashed line in Fig. 4.15. The resulting graphs are very similar to Fig. 4.16, and we can see more clearly that there is no power in the second waveguide at the zero-splitting point.

The maximum splitting observed in these experiments is in reasonable agreement with the theory. On X-2 (Fig. 4.16), we measure a maximum coupling rate $g_{\mathrm{ww}} = 3.8 \pm 0.1\,\mathrm{GHz}$,

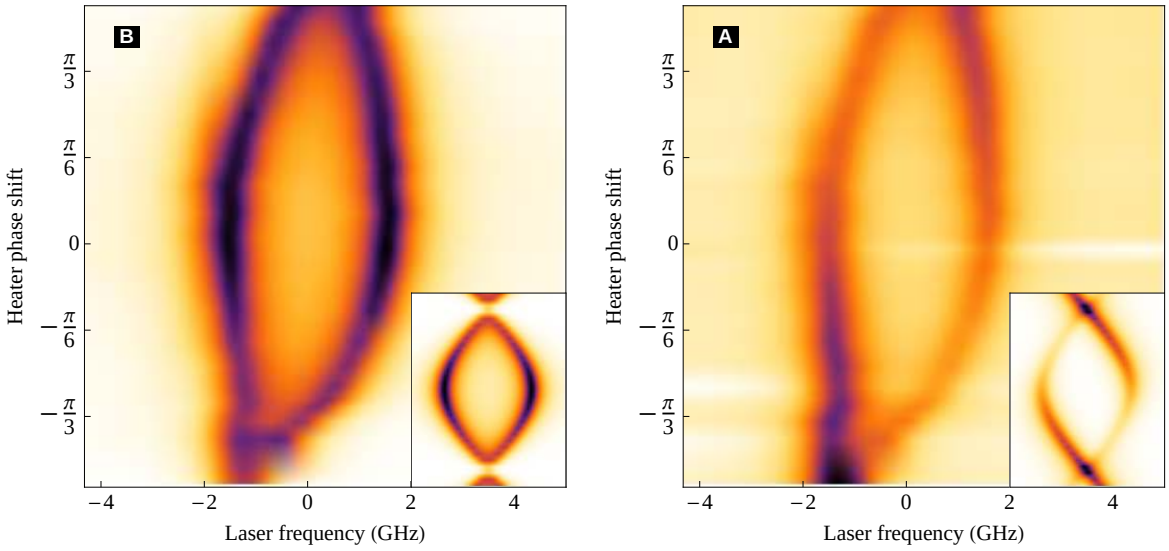

Figure 4.16: Tuning the coupling rate g_{ww} on coupler ORC8-X2. (B) and (A) are the transmission and reflection spectra respectively. The tuning is done by simultaneously varying the power across heaters H-A2 and H-B2 to produce a phase shift equal in magnitude (vertical axis), but opposite in sign, in the two heaters. The zero phase point has been arbitrarily set to coincide with the maximum coupling. Theoretical calculations are shown in inset.

while from a 30% coupling ratio we expected $3.9 \pm 0.2\,\mathrm{GHz}$. On X-1 (Figure 4.17) we measured $g_{\mathrm{ww}} = 3.5 \pm 0.1\,\mathrm{GHz}$, whereas we expected only $2.6 \pm 0.2\,\mathrm{GHz}$ ($C = 15\%$). It is quite possible that the initial measurement of the coupling ratio was off by as much as 50%: as we explained in Chapter 2, it is not as straightforward as it seems.

Conclusion

This concludes the main part of the present thesis. We have demonstrated the successful operation of two different types of coupled cavities: a waveguide resonator coupled to a free space microcavity, and two waveguide resonators with phase shifters for tuning. In both systems we reached the strong coupling regime, with a clearly resolved avoided crossing. The theory we developed, accounting for all sources of loss, reproduces quantitatively as well as qualitatively the main features of the data. On the other hand, our experimental results are rather anecdotal: the limited supply of waveguide chips as well as their fragility (notably the Nichrome heaters) did not allow for a systematic investigation of all the parameter space. Nevertheless, even on this limited basis, the evidence is certainly conclusive.

Having thus demonstrated the basic building blocks of the coupled-cavity QED implementation we devised in Chapter 1, building a full if modest array of up to a dozen cavities ought to be no more than an engineering problem, provided we can increase the finesse of the waveguide resonator. In our experiments, the very best exhibited a finesse

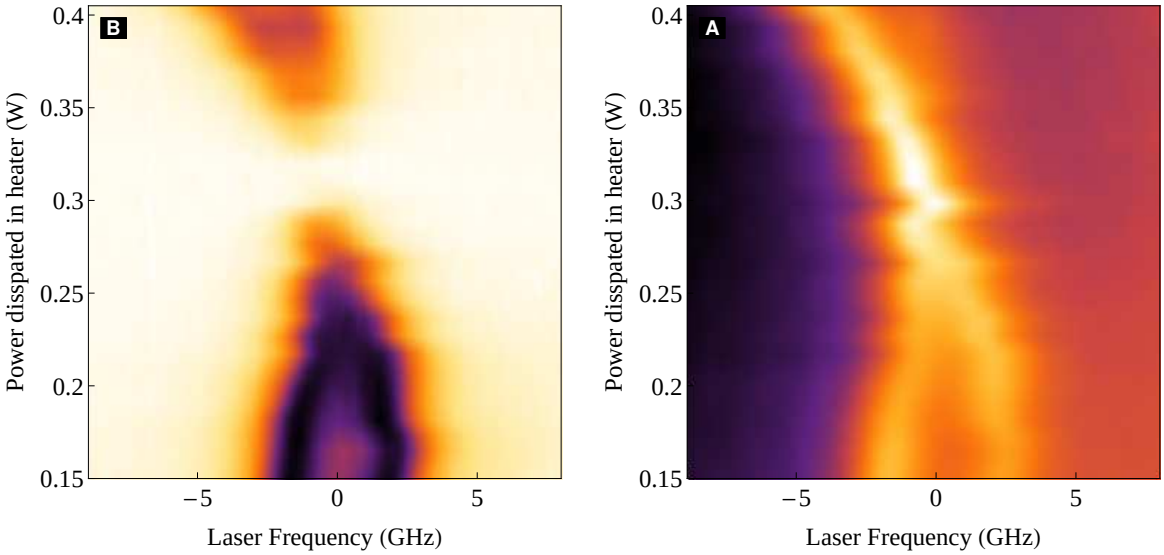

Figure 4.17: Tuning the coupling rate g_{ww} in cross-coupler ORC8-X1. The graphs show the transmission (waveguide B) and reflection (waveguide A) spectra as the the power in heater H-B1 is increased while the power in H-B2 is decreased by the same amount to tune the coupling rate, which corresponds to a change in the resonance's spectral splitting. The transmitted power vanishes as the coupling goes to zero near $P = 0.32\,\mathrm{W}$ (The heaters broke down before the power/phase shift calibration could be done).

of only 20 (16 if part of a coupler), whereas to resolve a spectrum with a dozen fringes will require a finesse of at least 100. And this does not even considers the more strenuous requirements of observing quantum effects.

We defer a fuller analysis of the loss problem to the final conclusion in Chapter 6. For now, it is time to look at the potential of Bragg gratings and their promise of an even tighter integration and better scalability.

Chapter 5

Bragg grating cavities

5.1 Introduction

When describing the UV-writing process (see Chapter 2.2.2), we mentioned the possibility of integrating Bragg gratings at essentially no cost. In fact the UV-writing system at ORC has such a capability, which we have not made use of yet. Since we do need many mirrors and cavities, the question naturally arises of whether we can find ways of exploiting Bragg gratings to reduce the cost, simplify fabrication, and reduce the footprint of our coupled cavities system. In particular the prospect of replacing the first waveguide mirror R_{w} (see Figure 4.1 on page 76) by a grating is very attractive. The waveguide chip would need to be mirror-coated on one end facet only, effectively halving the fabrication cost (currently, after the first mirror has been deposited, the chip needs to be taken out of the vacuum chamber, turned around and the vacuum needs to be restored). Moreover the waveguide cavities could be shortened by placing the gratings immediately after the couplers, maybe even into the S-bends, thus mitigating the propagation loss. The chip could also be made bigger and some elementary light processing (e.g. HBT interferometry) could be build into the same chip. The middle mirror R_{cw} might also be replaced by a grating, although mode-matching losses would likely reduce the microcavity finesse below acceptable levels; this would also significantly increase the cavity mode volume. Finally, removing the fragile end-facet DBRs would make handling the chip somewhat less risky.

For these reasons, we undertook to fabricate and study UV-written Bragg gratings. Although a mainstay of the telecom industry (where they are integrated, also by UV-writing, in optical fibres and operate at wavelengths around 1.5 μm), they are a rather new development in 780 nm waveguide chip technology. 1550 nm gratings had been made at the ORC for applications such a chemical and temperature sensing (113), but the 780 nm writing system was quite new when we embarked on this collaboration and we therefore entered uncharted territory: they were known to work, but had hardly been characterised at all and their actual reflectivities were unknown.

Beside potential applications in cavity QED, there is a lot of interest in integrated Bragg gratings at visible wavelength, mainly for biochemical sensing applications. "Lab-on-a-chip" devices, typically built around centimetre-scale glass or silicon-based planar devices, open a route to miniaturising the analysis of chemical and biological reactions thereby reducing the quantity of reactants required and the process time. One of the key requirements of this technology is an accurate method of monitoring environmental conditions and/or reaction progress. Fibre Bragg gratings can provide such monitoring by allowing part of the optical mode to interact with an analyte to measure small refractive index changes. It is therefore desirable to build Bragg gratings directly into the device, as has been demonstrated in (113) at 1550 nm by removing the cladding to allow a liquid analyte to interact with the optical mode of the waveguide. Such devices can display refractive index sensitivities better than 10^{-5} at this wavelength. However this sensitivity can be greatly compromised for water-based analytes due to the strong absorption of OH bonds, in excess of 30 dB/cm above 1400 nm (116), which reduces the finesse of grating and etalon-based sensors.

We also realised that cavities made of two Bragg gratings offered a quick and easy way of measuring not only the grating properties (reflectivity), but also those of the host waveguide (loss, refractive index). Most notably they offer a quick and easy way of measuring the gratings' group delay, a key property in the design of optical filters and delay lines and one which is usually measured by non-trivial means such as white light interferometry or pulse delay measurement.

This chapter is organised as follow: first we describe the simultaneous UV-writing of Bragg gratings by the stroboscopic method (the UV-writing process itself as been studied in a previous chapter). We then look into the theoretical modelling of the gratings in a coupled-mode formalism, with a detail analysis of the most important grating types we will have to deal with (uniform vs apodised and symmetric vs asymmetric). This is followed by a quick investigation of Bragg grating cavities, and how these allow one to measure the group delay through the indirect measurement of the cavity optical length. Turning next to the experiments, we describe the measurement methods and data analysis procedures before concluding on the grating properties.

Large sections of this chapter formed the basis for an article in *Optics Express* (141).

5.2 Direct UV-writing of Bragg gratings

The UV-writing set-up we have described in Chapter 2 (Fig. 2.3, page 36) is an earlier version of the actual set-up currently in use at ORC. Instead of having a single beam, the UV beam is split at a 50:50 beam splitter and the two beams are focused and overlapped in the core layer, where they produce an interference pattern with a periodicity, defined by the angle 2θ between the beams, of $\Lambda = \lambda_{\text{uv}}/2\sin\theta$, as shown in Fig. 5.1. Translation

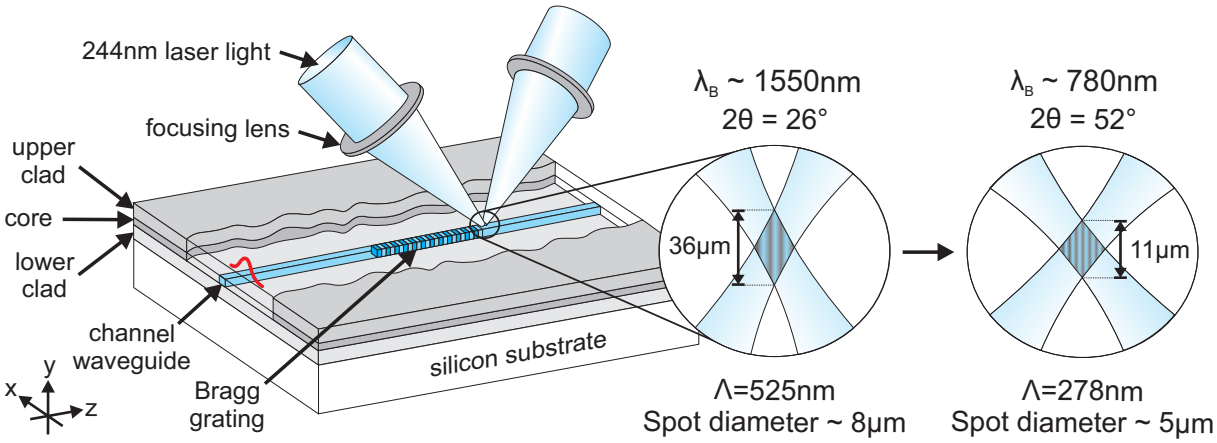

Figure 5.1: Direct Bragg grating UV writing scheme using 244 nm light. The Bragg wavelength Λ of the gratings is determined by the angle 2θ between the two interfering beams according to $\Lambda = 244\,\mathrm{nm}/(2\sin\theta)$). The two insets on the right indicate the main changes in the writing apparatus when moving from the traditional 1550nm regime to 780nm. Drawing courtesy of James Gates at the ORC.

of the sample under the CW laser spot averages out the interference pattern, resulting in a uniform waveguide indinstinguishable from those written with the single-beam apparatus. By contrast, if we switch the UV laser on and off by means of an acousto-optic modulator (AOM), with a period close to the time it takes for the sample to move through one period of the interference pattern, very much like a stroboscope, we will produce a periodic modulation of the refractive index, i.e. a Bragg grating.

We now move to a more detailed description of this process. The intensity profile of the intersecting beams (of waist w_z along the translation axis z) is given by

$$I(z) = I_0 \exp\left(-2\frac{z - z_0(t)}{w_z}\right)^2 \left[1 + \cos\left(\frac{2\pi(z - z_0(t))}{\Lambda}\right)\right]$$

as shown in the time domain in Fig. 5.2. $z_0(t) = vt$ is the position of the laser spot at time t, assuming that the sample is being translated at a constant speed v. At any fixed point A along the waveguide, located at z_A, the refractive index change will be proportional to the integrated intensity as A crosses the beams. This is called the fluence $F = \int_{t_{\mathrm{on}}} I(t)dt$. Since the laser is being modulated this integral is taken over the times when the laser is on, $t_{\mathrm{on}} = \{\Delta t_i\}$, and it decomposed as $F = \sum_i \int_{\Delta t_i} I(t)$. This integral has fortunately a somewhat decent analytical expression in term of the error function, which allows efficient numerical evaluation:

$$\int_{t_1}^{t_2} I(t)dt = \frac{w_z}{2v}\sqrt{\frac{\pi}{2}}\left\{\mathrm{erf}[b(t_2)] - \mathrm{erf}[b(t_1)] + e^{-a^2}\left(\mathrm{erf}[ia + b(t_2)] - \mathrm{erf}[ia + b(t_1)]\right)\right\}$$
$$\text{with} \quad a = \frac{\pi w_z}{\sqrt{2}\Lambda} \quad \text{and} \quad b(t) = \frac{\sqrt{2}}{w_z}(v.t - z) \tag{5.1}$$

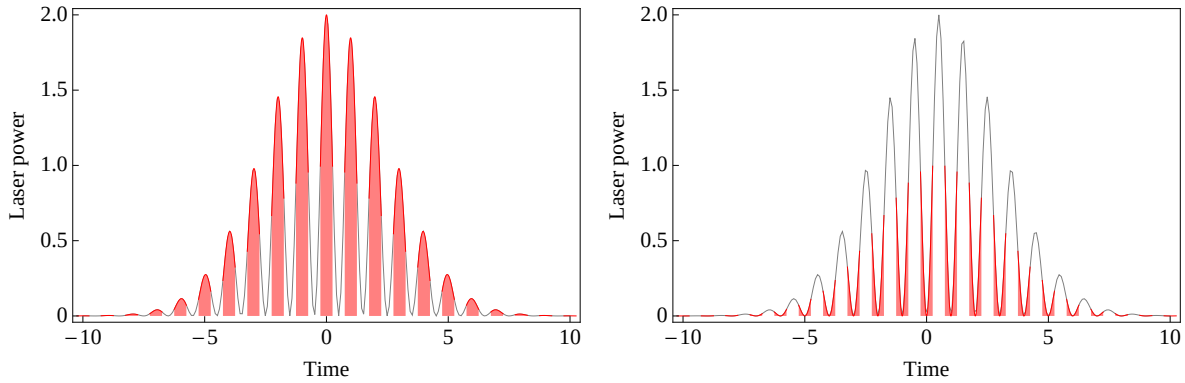

Figure 5.2: Principle of the "stroboscopic" method for direct UV-writing of Bragg gratings. The two plots represents the instantaneous UV power as a function of time at two points distant by half a grating period, as the sample is translated under the laser focus and its interference pattern (grey line). The laser is switched off outside of the regions marked in red, so that the total energy deposited at these points (the fluence) is proportional to the area under the red curve. The refractive index change in turn is proportional to the fluence, which reaches a maximum at $z = 0$ when the "on" times coincide with the maxima of the interference pattern, and a minimum a half interference period away, at $z = \Lambda/2$.

It is then straightforward to calculate the fluence along the grating. The relevant parameters are:

- **Duty cycle:** This is the fraction of a period during which the laser is on. A small duty cycle results in a sharper index profile and a stronger grating, as shown in Fig. 5.3.
- **Modulation period:** The grating period Λ (and therefore the Bragg wavelength λ_B) is set by the period Δt_{uv} of the laser on/off switching. Thus the grating design wavelength can be tuned, but only within some limits: If Δt_{uv} does not exactly match the period of the UV interference pattern, the stroboscopic overlap is not perfect and the fluence modulation contrast is smaller, resulting in weaker gratings (Fig. 5.4).
- **Writing speed:** The laser modulation has the unwanted effect of reducing the average fluence. In order to maintain F constant, the translation speed has to be reduced, otherwise both the effective index and the modal properties of the waveguide will not be matched to the other sections of the waveguide.

Fig. 5.2 illustrates the principle of the stroboscopic grating writing. It shows the laser power as a function of time at a fixed point along the grating; the red portions indicate when the laser is on. The fluence is proportional to the red-shaded area under the curve. In the left plot these coincide with the maxima of the interference pattern, so the fluence at that point is maximum. Contrast this with the plot on the right which represents a point located $\Lambda/2$ away from the previous one: the "on" periods now coincide with the minima of the interference pattern, resulting in a low fluence. Thus we produce a periodic

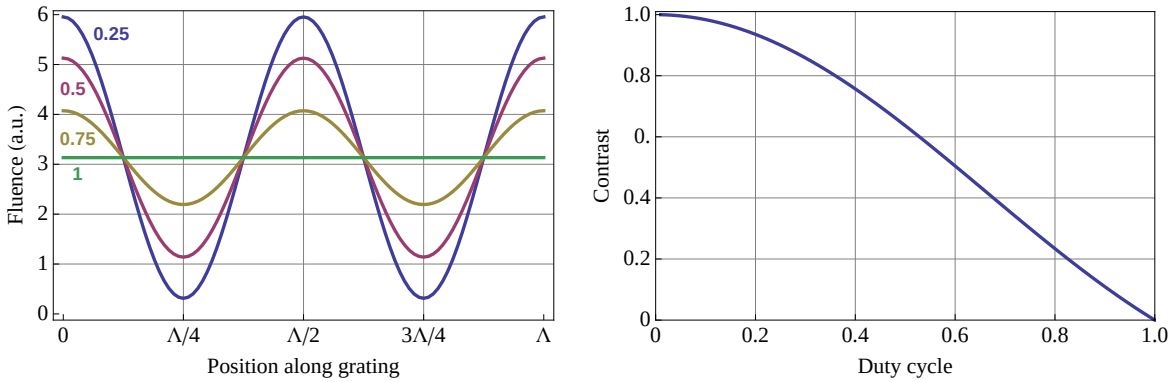

Figure 5.3: Left: fluence variation along two grating periods for various duty cycles. Right: fluence contrast as a function of duty cycle.

pattern of high and low index regions along the waveguide, with a smooth quasi-sinusoidal profile as shown on Fig. 5.3.

Duty cycle and grating contrast

Notice on the same figure how a smaller duty cycle results in a steeper grating profile as the "on" times get restricted to the most intense sections of the interference pattern. Ultimately the strength of the grating will be limited by the AOM resolution and by how long one is ready to wait for (since the writing speed is proportional to the duty cycle). The figure on the right shows the grating contrast as a function of the duty cycle. The latter can be adjusted along the grating, for example to obtain an apodised profile as discussed below.

Tuning the design wavelength

The grating Bragg wavelength λ_B can be adjusted via the UV laser modulation period. However this means that the on-off times and the UV interference pattern will no longer be fully synchronised, and the fluence contrast will be adversely affected. Figure 5.4 shows the fluence profile and contrast as the detuning is increased. It can be seens that this type of grating offers a relatively large range of working wavelength: indeed, if the UV interference pattern is optimised for operation at 780 nm, a detuning parameter of $d = 1.05$ (by which the contrast has dropped by about 25%) would correspond to a Bragg wavelength of 820 nm – meaning the UV-writing setup can produce gratings operating within a 80 nm band.

Effect of fluence saturation

So far we have assumed that the refractive index change is proportional to the fluence. However, as Figure 5.5(a) shows, the index change saturates at high fluence. This typically means that the upper envelop of the grating index profile, n_{high} will have a smaller contrast

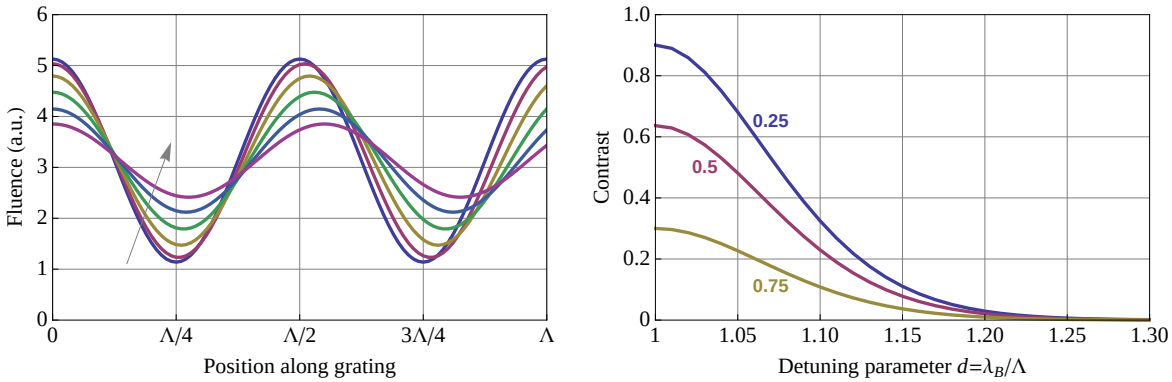

Figure 5.4: Effect of detuning on grating performance. Left, fluence profile along the waveguide as the detuning parameter $d = \lambda_B/\Lambda$ is increased from 1 to 1.1 (i.e. from 780 to 860nm along the grey arrow). The duty cycle is set to 0.5. Notice how the period of the fluence modulation increase, while the contrast drops. Right, the evolution of the contrast as a function of detuning, for three different duty cycles.

than the lower envelop n_{low}, as illustrated by Figure 5.5(b). As a result, the waveguide effective index $n_{\text{eff}} = (n_{\text{low}} + n_{\text{low}})/2$ will not be constant along the grating. The effect of this "DC-index change" (as opposed to the fast modulation with period Λ) is a spectral asymmetry in the grating response, which will be discussed in more details below (see Fig. 5.7).

Conclusion

In all of this discussion we have only considered ideal writing conditions. In practice the contrast of the gratings will be limited by that of the UV interference pattern and other experimental imperfections; but for our purpose of illustrating the grating production process this will suffice. Finally, it should be pointed out that an improved UV-writing set-up is currently being tested at the ORC, which offers significant improvements over the current stroboscopic method. The AOM switch is being replaced by a custom-made electro-optic modulator (EOM) inserted in one arm of the interferometer. Instead of modulating the laser, we can now translate the fringe pattern in sync with the sample by adjusting the EOM-induced phase shift. As a result there is no need to slow down the writing to accommodate a low duty cycle, and the relatively large loss introduced by the AOM (about 50%) is eliminated, making the writing process considerable faster. The grating contrast should also be significantly improved.

5.3 Modelling Bragg gratings

We have described how to integrate Bragg gratings to UV-written waveguides. We now move to a theoretical model of their operation. Before we start, let's clarify the terminology: by "Bragg gratings", we mean long (∼mm) Bragg structures made up of many

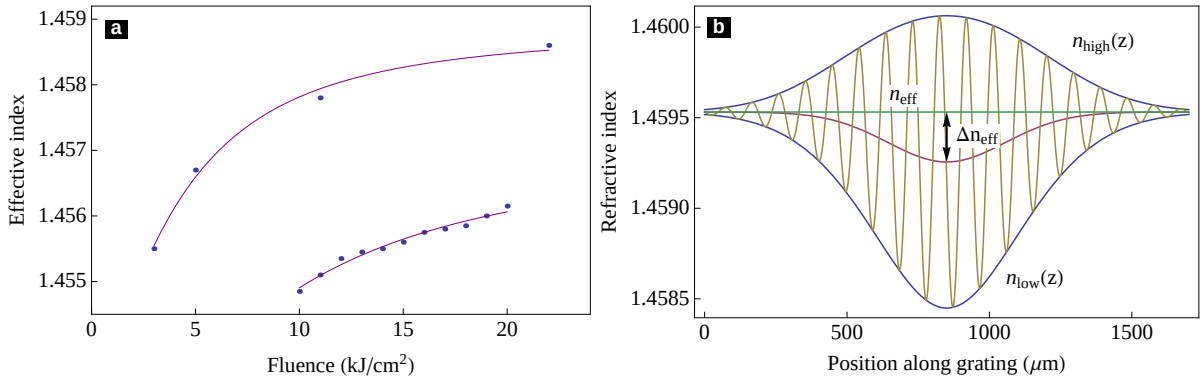

Figure 5.5: (a) Fluence saturation curves at 1560 nm for two different FHD samples, fitted with an empirical model. Our waveguides are made of the same FHD wafers that produced the lower curve, and where written with an average fluence of 14 kJ/cm^2. Data courtesy of Helen Rogers at the ORC. (b) Theoretical refractive index profile of a Bragg grating (yellow), assuming the lower fluence curve in (a), and with an additional 0.015 offset to take into account the different wavelengths (since our gratings are designed for 780 nm). The period has been exaggerated for the purpose of illustration. The high and low-index envelopes as well as the "DC" index change are also plotted: notice how the fluence saturation results in asymmetric gratings. See also Fig. 5.7 for the corresponding reflection spectrum.

periods, with a weak ($\sim 10^{-4}$), typically sinusoidal refractive index modulation such as those we just described, as opposed to "distributed Bragg reflectors" (DBR) or "Bragg mirrors" consisting of a small number (~ 10) of distinct high and low index layers with a very high index contrast (~ 1). Although obeying to the same principles, modelling their operation calls for a very different set of tools. DBRs are best described by looking at the Fresnel reflections at each interface and their coherent addition. This is easily done using one or another variant of the well-known transfer matrix method. In principle this could apply equally well to Bragg gratings, although the very large number of layers involved (especially for continuous index variations, where the index profile of each layer has to be discretised) makes it cumbersome and inefficient in practice. It is better to use a couple-mode framework in which a sinusoidal grating can be described entirely from the envelops and period of its refractive index modulation; for uniform gratings it is even possible to derive analytical solutions. In the following we will describe the coupled-mode formalism for Bragg gratings and highlight their most important features.

5.3.1 Couple-mode theory for Bragg gratings

We are only going to review the principles of coupled-mode theory in a very cursory fashion before focusing on its practical application to apodised and asymmetrical Bragg gratings. Erdogan (142) is an excellent and concise reference on this topic, although his notation is somewhat obscure.

A grating of length l is completely characterised by a small number of parameters, as shown in Figure 5.5: n_{eff} is the effective index of the host waveguide (as defined

in Chapter 2.1), $n_{\text{low}}(z)$ and $n_{\text{high}}(z)$ are the lower and higher index envelopes. These define the "DC" index change $\Delta n_{\text{eff}} = (n_{\text{high}}(z) + n_{\text{high}}(z))/2$ and the index contrast $\delta n = (n_{\text{high}}(z) - n_{\text{high}}(z))/2$. A grating with a design (aka Bragg) wavelength λ_B will have a spatial period $\Lambda = \lambda_d/2n_{\text{eff}}$.

If we only consider two identical counterpropagating modes inside the grating, with electric field amplitudes $A(z)$ and $B(z)$, and if we assume that the reflection of one into the other is the only process involved, we obtain a simple set of coupled equations (we refer the reader to Erdogan (142) for the complete derivation):

$$dA(z)/dt = +i(\sigma + i\gamma)A(z) + i\kappa B(z)e^{+2i\delta z} \tag{5.2}$$

$$dB(z)/dt = -i(\sigma + i\gamma)B(z) - i\kappa A(z)e^{-2i\delta z} \tag{5.3}$$

$\delta = 2\pi n_{\text{eff}}(1/\lambda_w - 1/\lambda_d)$ is the detuning between the design wavelength and the working wavelength λ_w; $\kappa = \pi\delta n/\lambda_w$ and $\sigma = 2\pi\Delta n_{\text{eff}}/\lambda_w$ are the cross- and self-coupling coefficient, respectively. We also introduced a loss parameter γ to describe absorption in the waveguide (the field amplitude then goes as $\exp(-\gamma z)$; we relate it to the other loss parameters introduced in previous chapters in Appendix C). It is straightforward to solve numerically this ordinary differential equation with initial conditions $A(0) = 1$ and $B(l) = 0$ for any working wavelength to establish the grating spectral properties. The reflectance, transmittance and phase shift upon reflection are then given by

$$R = |B(0)/A(0)|^2 \tag{5.4}$$

$$T = |A(l)/A(0)|^2 \tag{5.5}$$

$$\phi_r = Arg[B(0)/A(0)] \tag{5.6}$$

The grating strength depends on its length and index contrast. Analytical solutions do exist for the most simple case of uniform grating ($\delta n(z)$ =cst), in which case the reflectivity is given by $\tanh^2(\kappa L)$. But the grating is not completely conditioned by the κL product, because the phase-matching condition for constructive interferences inside the grating becomes more and more stringent as the number of layers increases, resulting in a smaller bandwidth even though the maximum reflectivity may be maintained constant by decreasing δn. Thus the bandwidth of a fibre or waveguide Bragg grating, with a typical δn of order $10^{-3} - 10^{-4}$, will be rather small at a few hundred GHz.

For non-uniform gratings there is no analytical solution to the coupled-mode equations, but they can be easily solved numerically. Thus we can calculate the spectral properties of gratings apodised in various ways (Gaussian, Cosine, Raised cosine and so on), of gratings with a phase defect (a half-wave layer at the centre that acts as a cavity), and even of arbitrarily complicated structures involving several gratings along a waveguide/fibre, like the grating cavities we will discuss shortly.

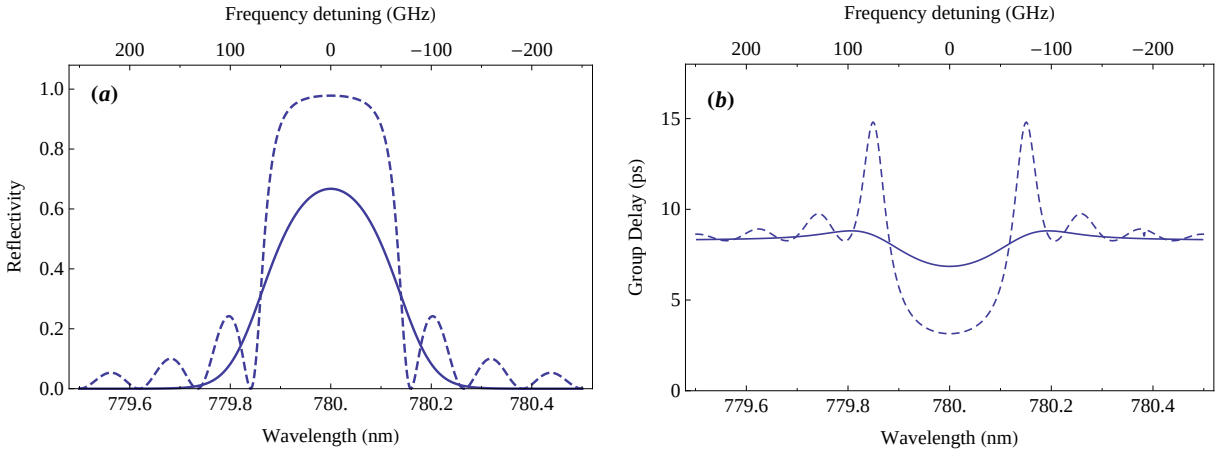

Figure 5.6: Reflectivity (left) and group delay (right) of uniform (dashed) and apodised gratings (solid line). The apodisation process removes undesirable ripples. The grating is 1.7 mm long and the index contrast is 0.00038.

5.3.2 Applications

Uniform and apodised gratings The simplest type of Bragg grating is that of a constant amplitude index modulation. This is the only type of grating to admit an analytical solution, and while they obviously offer the highest reflectivity for a given grating strength κL ($R_{\max} = \tanh^2(\kappa L)$), this comes at the cost of unwanted side lobes, as shown in Figure 5.6(a) (dashed line). These arise from the Fourier transform relationship between index profile and reflectivity (which is strictly true for weak gratings only). A widely used way of minimising the side lobes is through apodisation, that is by modulating the refractive index contrast with a smooth function, often a Gaussian. The spectral shape of the reflectivity function is then also close to a Gaussian (continuous line in Figure 5.6(a)).

Figure 5.6(b) shows the calculated reflection group delay $\tau_g = d\phi_r/d\omega$, which for simple grating structures is just a measure of the field penetration inside the grating. τ_g is smaller where the reflectivity is maximum, because the field gets reflected earlier. There too, apodisation reduces the amplitude of the ripples. All of the gratings used in this work were Gaussian-apodised (with a $1/e$ half-width equal to a quarter of the grating length).

Symmetrical and asymmetrical gratings We have discussed the effect of fluence saturation on the grating index profile. The main effect is that the upper and lower envelopes are no longer symmetrical. Figure 5.7 illustrates how this affects the reflectivity profile and group delay in the case of an apodised grating. The reflectivity acquires a pronounced asymmetry and small sidelobes on the short wavelength side (near 779.8 nm). This is because the DC-index change causes shorter wavelengths to be closer to the design wavelength in the wings of the gratings than in the centre, which results in a small

enhancement of the reflectivity together with the formation of an optical cavity. The latter causes the zero in the reflection spectrum. It is believed that the asymmetry often observed in the experiments we will describe later result mostly from this DC-index effect.

The overall shift in the Bragg wavelength, $\Delta\lambda_B$, results from the the lower effective index at the centre of the grating and is given by $\Delta\lambda_B = \frac{\Delta n_{\text{eff}}}{n_{\text{eff}}}\lambda_B$

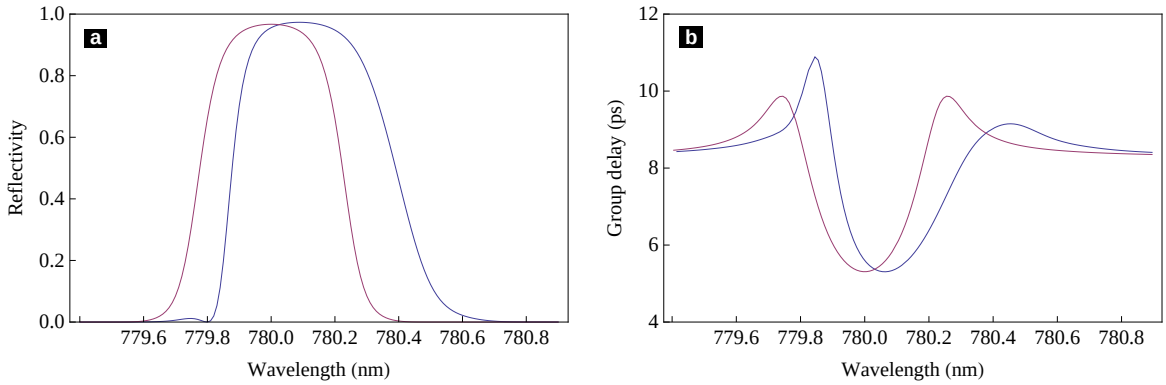

Figure 5.7: Typical reflectivity and group delay of a grating with and without dc-index. The presence of a dc-index induces an asymmetry of the reflection spectra as well as additional ripples on the blue side and general spectral shift. The asymmetric grating profile is calculated from the lower fluence curve of Figure 5.5(a).

Bragg cavities and group delay The group delay (ie the delay experienced by the reflected wave upon reflection) is given by

$$\tau_g = \frac{d\phi}{d\omega} = -\frac{\lambda^2}{2\pi c}\frac{d\phi}{d\lambda} \tag{5.7}$$

and is typically measured in picoseconds. The group delay can be related to an *effective penetration length* L_g of the incoming wave into the grating through $\tau_g = 2L_g n_{\text{eff}}/c$ (n_{eff} is the effective index of the host waveguide). To validate this interpretation, one can imagine a cavity made of two identical gratings, separated by a distance L, and measure its free spectral range FSR $= c/2n_{\text{eff}}L_{\text{eff}}$. The effective length of the cavity should be $L_{\text{eff}} = L + 2L_g$.

We simulated such a cavity with two apodised gratings of length 1.7 mm and index contrast 8×10^{-4}, with a 14.3 mm gap (i.e. the maxima of their index envelops are separated by 16 mm). The resulting spectrum in Figure 5.8(a) exhibits the usual Fabry-Perot comb structure (the gap in this figure has been reduced to 4 mm for clarity). From the spectrum, we extracted the frequency of about 60 resonance fringes to calculate the free spectral range and, through the equations given above, we estimated the grating group delay τ_g (plotted as red dots in (Figure 5.8(a)). These can be compared with the group delay calculated directly for a single grating (blue line), and the agreement validates our initial interpretation of the group delay in term of penetration depth.

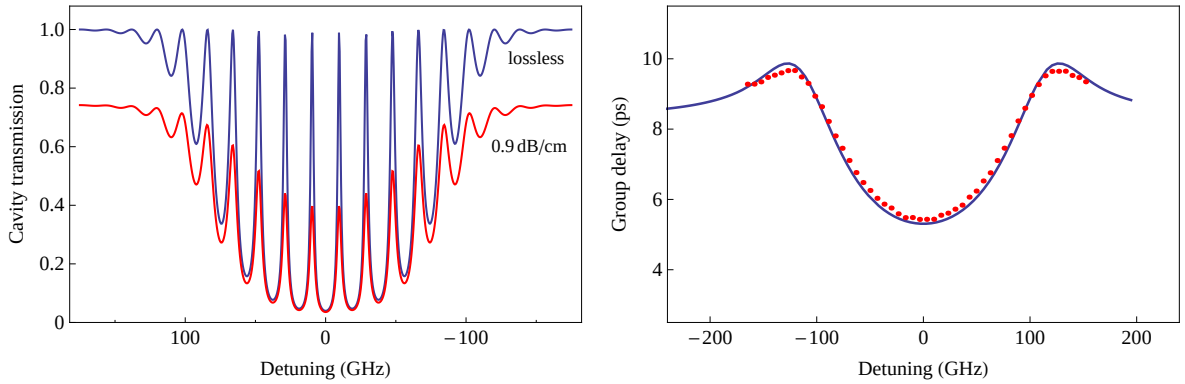

Figure 5.8: (a) Calculated transmission spectrum of a cavity made up of two Gaussian-apodised gratings separated by 4 mm. The blue curves shows the lossless case whereas the red curve includes realistic propagation losses of 0.9 dB/cm. (b) Group delay as calculated from coupled-modes theory (solid line) and derived from a cavity spectrum similar to that on the left.

Impact of propagation loss Finally, a few words about losses are in order. The grating cavity spectrum we just discussed (Figure 5.8(a), blue curve) did not include any losses. The red curve in the same plot does include a propagation loss of 0.9 dB/cm, typical of out worse waveguides. As can be seen by looking at the edges of the grating stop-band, this corresponds to a single-pass loss of 25% in this 4 mm-long cavity. The finesse drops considerably, and the amplitude of the resonance peaks (or, if we were looking at the reflection spectrum, the visibility) also decreases. This is a consequence of the larger photon lifetime at higher reflectivities, which increase the propability of absorption due to propagation loss. Another factor in the reduced visibility is the mismatch effectively induced between the two mirror. All of this can be easily understood from the detailed discussion on lossy cavities in Chapter 3.2.

5.4 Bragg grating cavities: experiments

Having described the fabrication, modelling and properties of Bragg gratings, we now turn to the experimental characterisation of actual gratings. All the following experiments were performed on a single waveguide chip (codename: ORC3) containing six sets of five waveguides with an integrated grating cavity. The grating length varies between the sets (denoted A, B,..., F) from 1.3, 1.4,..., to 1.8 mm, while within each set the gap between the gratings varies between 4 and 16 mm. All waveguide were written at the same fluence of 14 mJ/cm^2.

We undertook two sets of experiments. First, we measured finesse, transmission and FSR of all 30 cavities at the Bragg wavelength, from which we estimated the mirror reflectivities, penetration lengths, and the waveguide propagation loss. We then focused on set E (1.7 mm-long gratings) and measured the same quantities over a large bandwidth, mainly in order to demonstrate the group delay reconstruction from FSR discussed earlier.

But before we can look at these results, we need to describe the experimental and data analysis methods.

5.4.1 Apparatus

The experimental setup is illustrated in Figure 5.9. Light from two 780 nm distributed feedback (DFB) laser diodes (Eagleyard Photonics) is coupled to polarisation-maintaining fibres, and then combined at a beam splitter so that either laser can be used to probe the waveguides. Typically one laser is tuned far from the Bragg frequency and is used to measure the transmission through the waveguides in the "absence" of gratings. The second laser is tuned over up to 100 GHz by scanning the drive current of the laser diode, and beyond that by changing its temperature. Half- and quarter-wave plates adjust the polarisation before the light is coupled into the single-mode input fibre. The exit of this fibre is cleaved and aligned with a waveguide using a precision three-axis flexure stage. A second such stage aligns the output fibre, which is then fed to a photodiode. Undesirable etalon effects between the fibres and the waveguide chip are avoided by using index-matching oil.

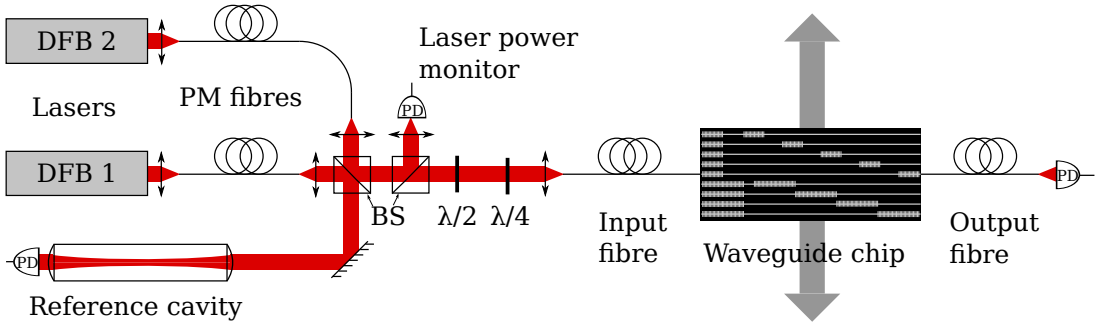

Figure 5.9: Schematic representation of the experimental setup. On the chip, a typical set of cavities with different grating and cavity lengths.

A second beam splitter sends a fraction of the light to a photodiode used as a normalisation to compensate for variations of the laser power during a frequency scan. Finally, part of the scanning laser is diverted to a reference cavity consisting of a 5 cm long block of glass, with two reflection-coated lenses glued to the ends of the block. This cavity sits in a heavy aluminium housing that provides thermal inertia. A cavity finesse of 30 is measured, in agreement with the nominal mirror reflectivity of 90%. The free spectral range is measured to be $\mathrm{FSR}_{\mathrm{ref}} = 1.905(5)\,\mathrm{GHz}$. During a laser scan we monitor the fringes of light transmitted by this cavity to obtain a calibrated relative frequency scale.

5.4.2 Measurement method

Broadband frequency scanning

As we have seen, the bandwidth of typical UV-written Bragg gratings is of the order of 200 GHz. The mode-hop free current-tuning range of the laser diode however is only about 100 GHz. Therefore to scan the laser frequency over the required 500 GHz range, we use a combination of current and temperature tuning. While scanning the current rapidly (20 Hz) over a few FSRs, we change the set-point of the laser temperature controller. As the temperature, and therefore the mean laser frequency, slowly drifts over several minutes, we accumulate a collection of transmission spectra (also called frames) for both the reference and waveguide cavities. The accurately known (and constant) FSR of the reference cavity is used to normalise the frequency, i.e. to convert the time axis of each frame to optical frequency, as detailed in the next section.

Because the temperature-induced frequency shift is much slower than the current one (as illustrated in Figure 5.10), there is a large overlap between successive frames. In particular, if the mean frequency shift from one frame to the next is smaller than the FSR, it is possible to keep track of the phase and frequency over the entire collection of frame, and thus to reconstruct the whole spectrum.

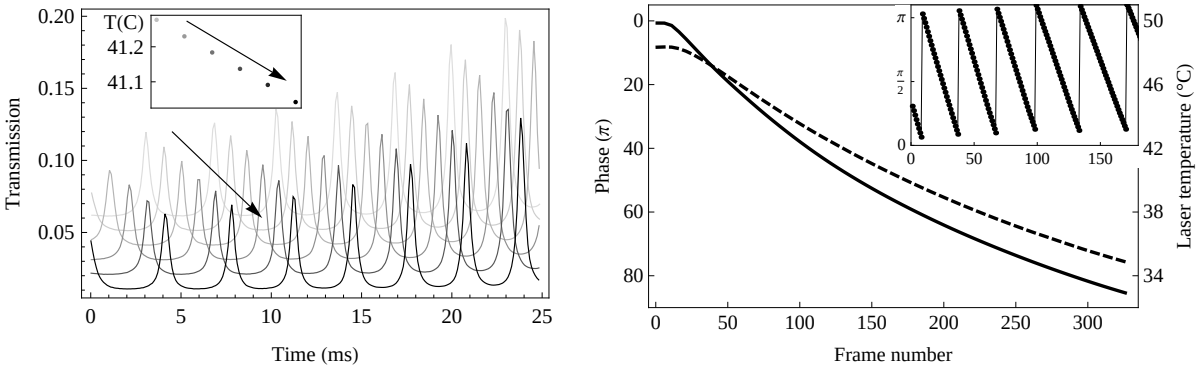

Figure 5.10: Left: Six successive transmission spectra (frames) for the 16 mm waveguide cavity. The large overlap makes it possible to keep track of the slowly varying initial frequency. Inset: the laser diode temperature (Celsius) at which each frame was taken. The arrow indicates the order in which the frames were taken. Right: Typical evolution of the laser diode temperature (dashed line) and reconstructed phase of the grating cavity spectra (solid line) during one measurement. The total scan amplitude is about 90 times the cavity FSR, or 500 GHz over 12 °C. The inset shows the phase evolution before unwrapping.

Frequency normalisation

The relationship between the laser current and frequency is slightly non-linear, as can be seen in Figure 5.11 and 5.10 where the FSR of the reference and waveguide cavities is clearly larger at the beginning of the frame. Therefore, before we can fit the waveguide cavity spectrum with an Airy function to extract information such as finesse, free spectral

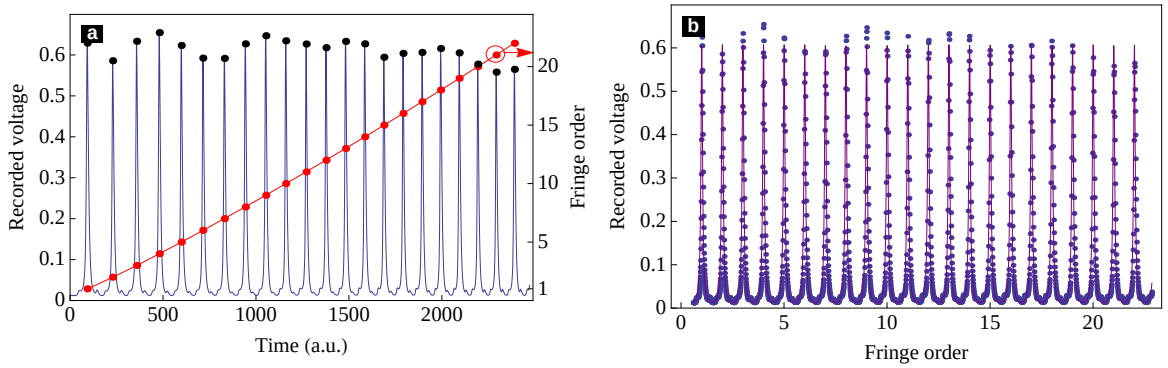

Figure 5.11: Frequency normalisation. Left, the reference cavity spectrum as acquired. The red dots indicates the position of each resonance fringe, which is very non-linear in time/drive current, fitted with Eg. 5.8. Left: the same spectrum after normalisation, fitted with an Airy function. The frequency is now expressed in terms of the reference cavity FSR, 1.905 GHz.

range etc., we have to reconstruct the frequency scale. This is the role of the reference cavity described previously.

The normalisation process is summarised in Figure 5.11. A typical frame (blue curve in figure (a)) consists of 2000-2500 samples acquired via the NI-DAQmx card at a rate of up to 400 kS/s. For each of the $p = \{1, ..., n\}$ fringes in the spectrum we recover the timestamp of the sample of maximum amplitude (black dots). We then plot the order number as a function of time, $p(t)$ (red dots) and fit it with an empirical function of the form

$$p(t) = a + b\sqrt{t} + ct + dt^2. \tag{5.8}$$

Although it is possible to improve the accuracy of the fringe timestamping, for example by fitting a quadratic polynomial around each resonance, this has not proved necessary, partly because any jitter in the determination of fringe position will be smoothed out during the fitting of Eq. (5.8). The last step is simply to replot the spectrum against $p(t)$ (Fig. 5.11(b)). The frequency will be in units of the reference cavity FSR, and conversion to physical units is immediate. Figure 5.11(b) shows the initial reference cavity spectrum replotted against the normalised frequency, and fitted with an Airy function whose FSR has been set to 1. The finesse is 30, as expected from the 90% reflectivity of the lens coating, and the fit is excellent.

Fitting the spectra

Once all the frames within an experimental run have been individually normalised, they can be fitted with the usual Airy function $A(f)$ of Eq. (3.7). Since the absolute frequency of the fringes in not known, we include a frequency offset $f_0 = \psi \times \text{FSR}$; we also need a

DC offset b:

$$A(f) = b + \frac{I}{1 + F \sin^2\left[\pi\left(\frac{f}{\mathrm{FSR}} + \psi\right)\right]}. \tag{5.9}$$

The reflectivity of the gratings can vary substantially within one frame, so that the finesse F, free spectral range FSR and amplitude I are not constant. Therefore we let them vary linearly or quadratically, on a case-by-case basis, e.g. $F(f) = F_0 + F_1 \times f + F_2 \times f^2$. F_0 is then taken as "the" frame finesse, and so on.

With such a rapidly varying function up to a dozen free parameters, the proper convergence of the fitting algorithm depends on good initial guesses for these parameters. With hundreds of frames in each experimental run, manual fine-tuning is out of the question, so we use a semi-automated method as follows.

A single initial guess for the finesse F_{init} and free spectral range $\mathrm{FSR}_{\mathrm{init}}$ is usually sufficient for a whole set of frames, and the amplitude I_{init} can be easily estimated for each frame. Much more critical is the phase, for which we use the following estimate:

$$\psi_{\mathrm{init}} = \frac{f_m \mod \mathrm{FSR}_{\mathrm{init}}}{\mathrm{FSR}_{\mathrm{init}}} \tag{5.10}$$

where f_m is the frequency of the resonance with the highest amplitude in the frame, and mod denotes the modulo operation. Even with such precautions it is not uncommon for the fit to fail in a significant proportion of frames, as exemplified by the scattered blue points in Figure 5.12(a,b,c).

To improve the fits, we devised a feedback method based on the idea that the parameters cannot differ much from one frame to the next, and therefore when the fit fails it can be improved by updating the starting points using values from nearby, successful fits. We do so by cleaning up the parameters of the first fit attempt with a combination of thresholding and a moving average. The cleaned-up parameters (Figure 5.12(a,b,c), purple points) are then fed back to the fit routine as new initial guesses, and the second fit is usually fairly successful, as shown in Figure 5.12(d). A few frames may still not be fitted properly and can be fine-tuned by hand, or even just ignored. Occasionally repeating the feedback process one more time, including by changing the expansion order of F, I or FSR, can yields significant improvements. This feedback method is semi-automated in the sense that no fine-tuning of individual frames is required, but that manual intervention is necessary to properly clean up the initial fit parameters.

Once we have determined all the parameters relevant to single frames, we need to attribute a single frequency f_i to the frames in order to reconstruct the grating spectrum $F(f_i)$, $\mathrm{FSR}(f_i)$ etc. From our definition of the frame phase ψ_i, we have the following the relation between two consecutive frames:

$$f_{i+1} - f_i = \mathrm{FSR}_{\mathrm{ref}} \times \mathrm{FSR}_i \times (\psi_{i+1} - \psi_i). \tag{5.11}$$

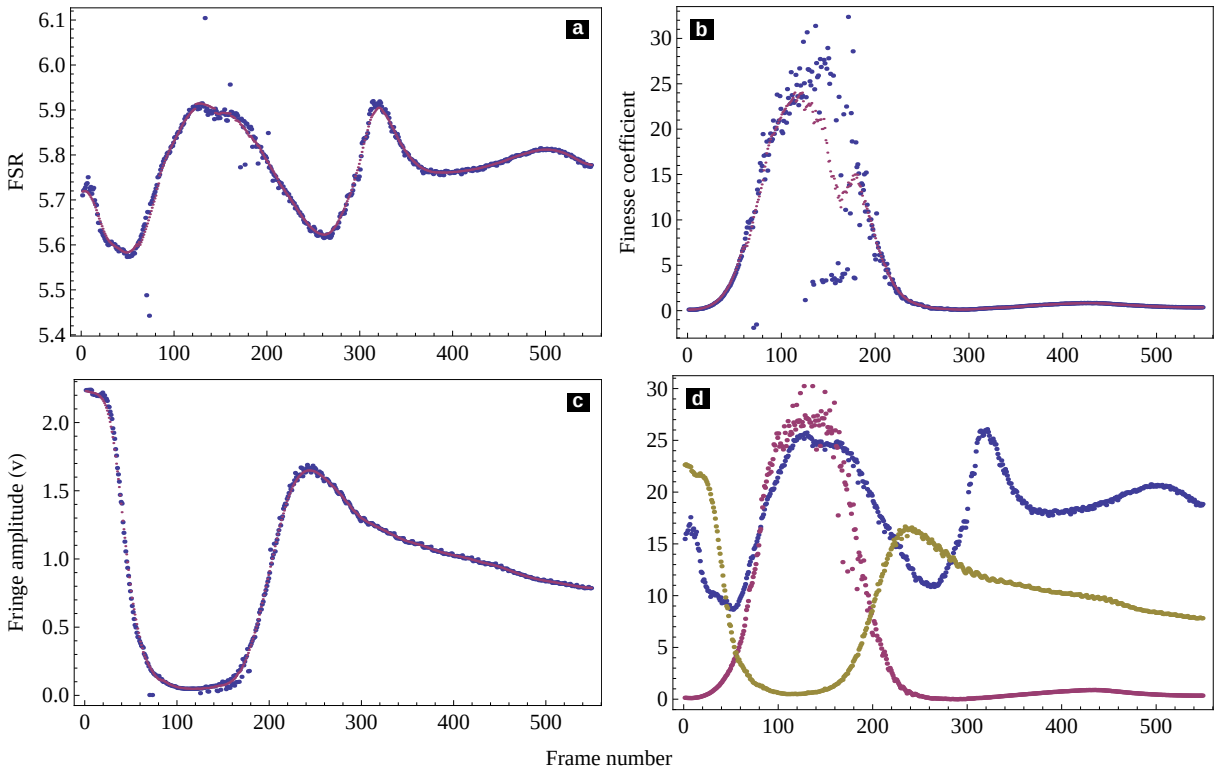

Figure 5.12: Adaptive fitting. Plots (a), (b) and (c) show the free spectral range, finesse and fringe amplitude as determined by the initial fit (blue dots) and after the clean-up procedure (purple dots). (d) shows the same parameters at the end of the adaptive fitting process; there are no outliers left.

To reconstruct f_i fully, we need to unwrap the phase, as illustrated in Figure 5.10(b). This simply involves parsing the list of $\{\psi_i\}$ and adding either 1 or -1 (conditioned on the phase increasing or decreasing) to all subsequent ψ_i whenever ψ reaches 1 and is reset to 0. Figure 5.10(b) shows a typical example of the phase evolution across a few hundred frames, before and after unwrapping, as the laser temperature changes by about 12°C.

Free spectral range noise improvement

One of the most important feature of this convoluted data acquisition process is that it allows the determination of the laser frequency throughout a large temperature scan without the need for a wavemeter. One can judge the surfeit of data to be positive or negative depending on the perspective. In any case it is fair to argue that much of it is redundant: the FSR as much as the finesse are after all only defined with respect to a particular resonance, or a pair thereof. One is therefore tempted to use this redundancy to obtain a single, averaged value of the cavity parameters, associated with each resonance.

A more practical motivation for doing so is the presence of oscillatory artefact with a period identical to the waveguide cavity FSR, particularly on the FSR itself, as exemplified in Figure 5.13. These artefacts are a rather unavoidable side-effect of the fitting procedure: it is not hard to see why the sudden entrance of a new fringe in a frame, or the exit of

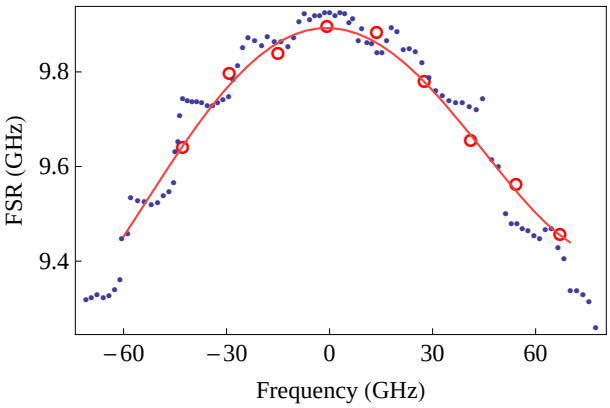

Figure 5.13: A FSR spectrum reconstructed through the fitting method (blue dots), exhibiting a typical staircase structure associated with new resonances entering or exiting the frames. The red circles show the result of the data reduction method (see text for details), where each resoance order has been attributed a single FSR value. The red curve is just a guide for the eye.

another one, could and does cause jumps in the fitted parameters. We have to alleviate this problem in one way or another before we can focus our attention on the FSR, and use it to reconstruct the group delay.

The starting point of the method we are about to describe is the observation that the frame phase obtained through the fitting process, ψ_i, is immune to these artefacts. We can therefore take it for granted. We can also take the average of all the fits FSR as a crude global estimate of the FSR, $\mathrm{FSR}_{\mathrm{init}}$. We then reanalyse the raw data by fitting a quadratic polynomial around the tip of every fringe, in order to extract accurately the *relative frequency* $f_{i,j}$ of all fringe orders j in all frames i. The same fringe of order j will appear across several frames with a different relative frequency $f_{i,j}$, but the *absolute frequency* of that particular order, ν_j, should be the same. Thus we produce Figure 5.14(a), which plot the estimated absolute frequency

$$\nu_j = f_{i,j} - \mathrm{FSR}_{\mathrm{init}} \times \psi_i \tag{5.12}$$

of the nearly 4000 fringes in this data set (on average 8 fringes in each of 500 frame). We then produce Figure 5.14(b) by gathering togeter all fringes belonging to the same order. Because we used only an estimate of the FSR, the ν_j are not all identical, as in the red inset. But to get the "true" FSR, all we have to do is adjust $\mathrm{FSR}_{\mathrm{init}}$ until they form an horizontal line (green inset). Mathematically we express this as

$$f_{i,j} = \nu_j + \mathrm{FSR}_j \times \psi_i. \tag{5.13}$$

This set of m linear equations forms a trivial linear least squares problem and can easily be solved for ν_j and FSR_j to obtain exactly what we were looking for: a single, "averaged" (in this case through the least square fit) FSR associated with each resonance order.

Figure 5.13 compares the FSR before and after application of this method: notice how the artefact is all but gone.

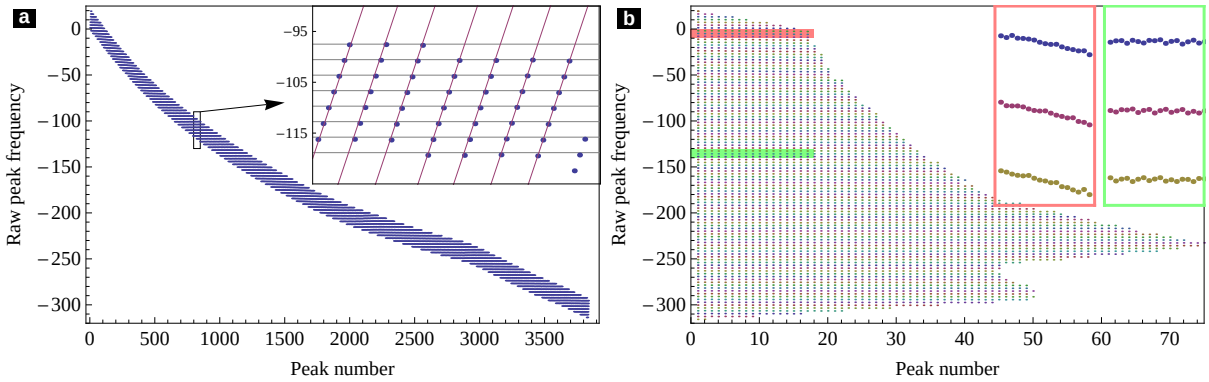

Figure 5.14: Illustration of the improved FSR extraction method. Left, the absolute frequency of all resonance peaks from all frames, $f_{i,j} = f_i + \mathrm{FSR}_{\mathrm{init}} \times \psi_i$. Notice the pattern shown in the inset: the purple, diagonal lines indicate resonances from a single frame, whereas the grey horizontal lines follow a single resonance across consecutive frames. Grouping together the resonances of the same order across many frames is then straightforward, as illustrated in the figure on the right. This shows that an inaccurate value of $\mathrm{FSR}_{\mathrm{init}}$ results in a slow drift of the calculated n^{th} fringe frequency. A "time-averaged" FSR can thus be determined by minimising this drift.

5.4.3 Experimental results I

A typical, large-range set of data is shown in Figure 5.15 for a 1.7 mm-long grating cavity with 16 mm spacing (E1). Graph (a) shows, on a logarithmic scale, a few frames extracted from a total of about 400 to recontruct the spectrum over the whole grating stopband, while in (b) we plot the finesse, FSR and fringe amplitude. All three parameters behave according to the patterns discussed earlier. The larger FSR at the centre of the stop band indicates a smaller penetration length where the reflectivity (and the finesse F) is higher, while the dramatic drop in fringe amplitude I (almost two orders of magnitude) betrays the presence of large propagation loss. The clear asymmetry, with a sharper finesse drop at shorter wavelengths, is most likely a consequence of the fluence saturation and resulting DC component of the index modulation.

We then measured these three parameters (F, I, FSR) at the centre of the stop band for all 30 waveguides. The results are summarised by sets of identical grating lengths in Figure 5.16 for F and I, for the two normal waveguide polarisations. The propagation loss and mirror reflectivities calculated according to Eq. (3.24) are also shown.

To help with the interpretation we have plotted the theoretical finesse and transmission (grey lines in set F, based on typical parameters $R = 0.84$ and 0.9 dB/cm loss), as well as the theoretical reflectivity (red dots in the last column, assuming an index contrast $\Delta n = 6.4 \times 10^{-4}$). The purpose of these indicator is to give an idea of the expected trends:

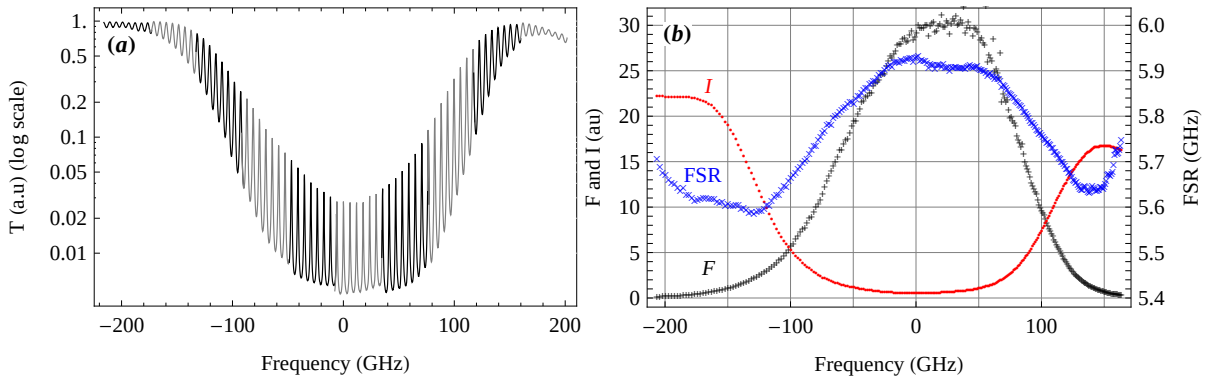

Figure 5.15: Data taken using a 16 mm long waveguide cavity with 1.7 mm-long gratings. (a) Log of intensity transmitted by cavity, reconstituted from 13 frames (out of 400), with gray/black parts denoting different frames. In the wings the transmission is close to 100%, at the centre it oscillates between 0.5% and 3%. (b) Black pluses show the coefficient of finesse F, blue crosses plot the free spectral range FSR, and red dots indicate the maximum transmitted intensity.

they are not intended to fit the data. Important deviations from these expectations are clearly visible. First, for each set of waveguides, we expect the total propagation loss to increase with the cavity length, and therefore the finesse and transmission should also fall. This is indeed what we observe for the transmission, but when it comes to the finesse this trend is often violated. Set A is most peculiar in that regard; moreover it shows a finesse significantly higher than the other sets, although the gratings are shorter and the reflectivity ought to be smaller. The mirror reflectivity should be constant within each set, and increasing with grating length. Neither seems to be the case, with R varying considerably between 70 and 90%, and the loss between 0.5 and 1.5 dB/cm. The one consistent trend throughout the data is that the propagation loss is slightly lower for one of the two polarisations.

When using Eq. (3.24) to calculate loss and reflectivity, we did not include the mirror loss A. However, introducing this parameter, assumed to be constant for all gratings, does nothing to alleviate the inconsistencies, so we are left with the conclusion that they are simply a natural product of the writing process. This lack of reproducibility, and maybe more importantly the fact that the reflectivity seems limited to about 90%, make the prospect of using these gratings in a cavity QED system highly unrealistic.

Let's now turn our attention to the free spectral range (Figure 5.17). Since the FSR is inversely proportional to the effective cavity length L_{eff}, a plot of FSR^{-1} versus L yields an intercept on the abscissa of $-2L_g$, which gives τ_g. Figure 5.17(a) shows the plots for two series of cavities, one with gratings 1.4 mm long and the other 1.8 mmS. The six values of τ_g derived from these and the other four sets of cavities are plotted in Fig. 5.17(b) as black dots with error bars. As indicated by the dashed line, τ_g increases linearly with the physical length of the grating. In part, this reflects the fact that, because of the

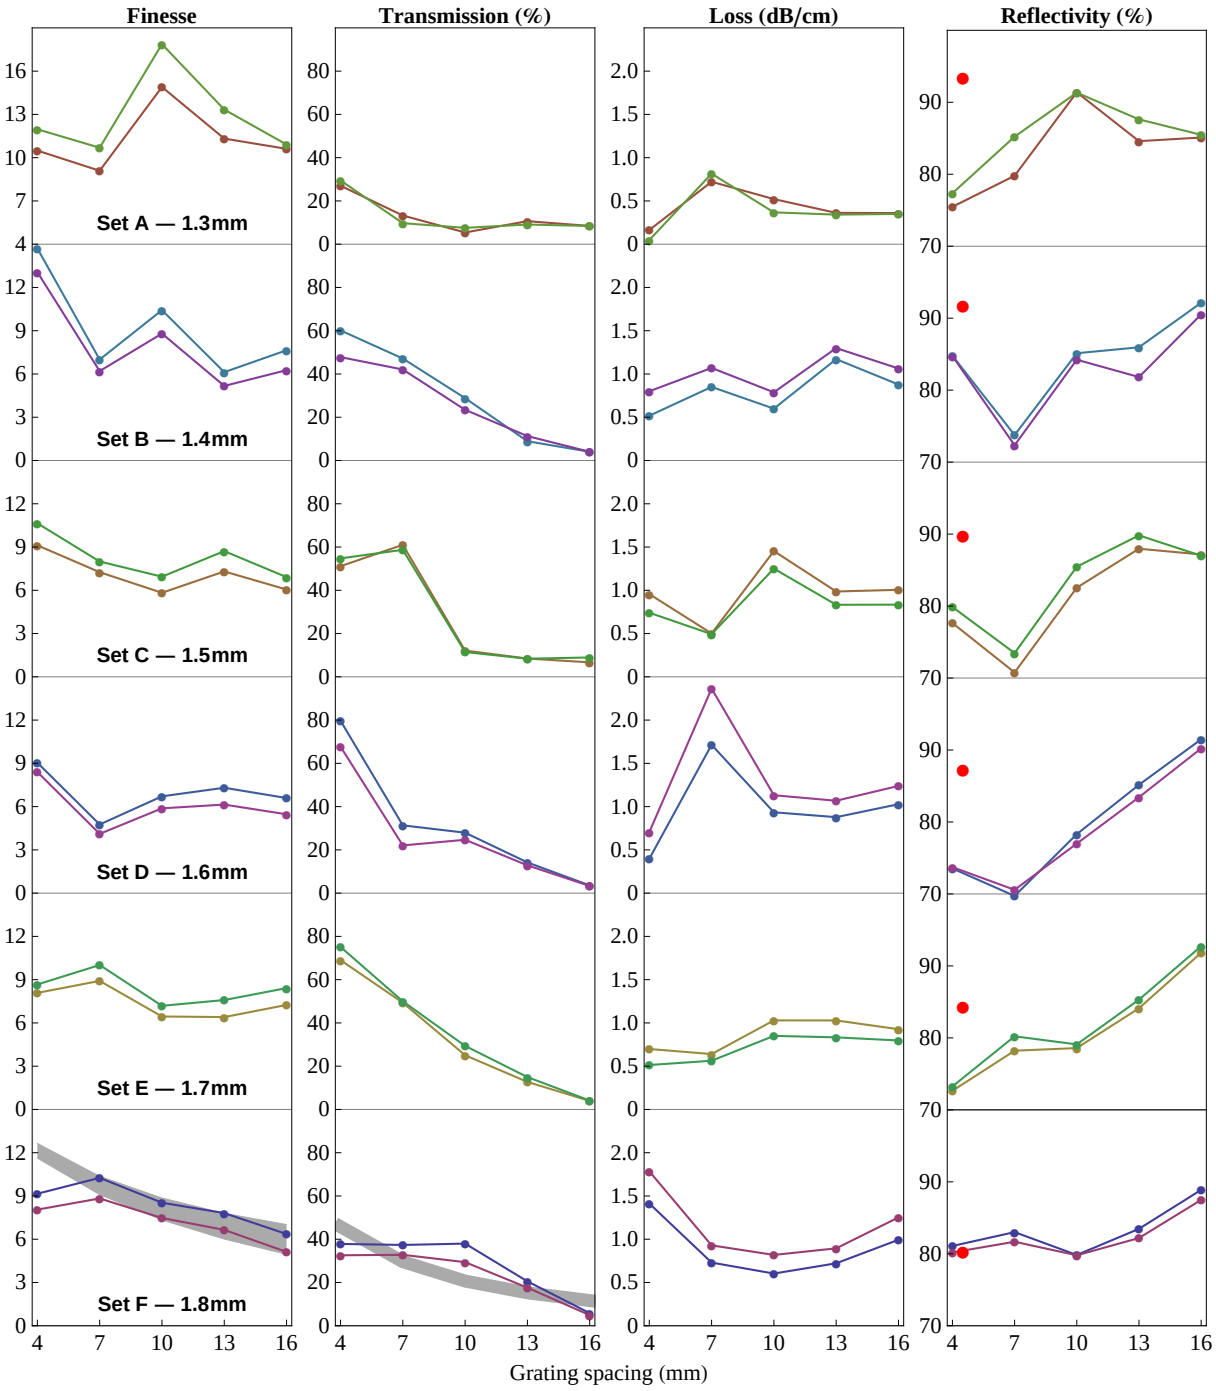

Figure 5.16: Summary of measurements on chip ORC3, grating sets A to . From top to bottom, the grating lengths are 1.3, 1.5,...,1.8 mm respectively. First two columns: Measured finesse and transmission at the centre of the grating's linewidth. Last two columns: Propagation loss and mirror reflectivity as calculated from the data on the left, according to Eq. (3.24). On each plot the two sets of data represent the two orthogonal polarisations. The grey lines on the last line show the theoretical behaviour of F and T based on $R = 0.84$ and 0.9 dB/cm propagation loss; the red dots in the last column are the theoretical mirror reflectivities based on $\Delta n = 6.4 \times 10^{-4}$. They have been added to illustrates the expected trend: the lack of consistency reflects large uncertainties in the UV-writing process.

apodising function, the field must penetrate more deeply into longer gratings before the index contrast is high enough to produce strong reflection. The upper solid red line in Fig. 5.17(b) show the group delay given by coupled mode theory for an index contrast of 3.9×10^{-4}. By contrast, the results obtained using shorter gratings correspond to Δn as large as 6.3×10^{-4}, the lower red line. We believe that the decrease of Δn with grating length is in fact a decrease of index contrast with the time of writing, the gratings having been written in order of increasing length. The waveguide chip is loaded with hydrogen prior to writing in order to increase the photosensitivity, and it is known that enough hydrogen outgasses over the several hours of writing to produce an effect of this size.

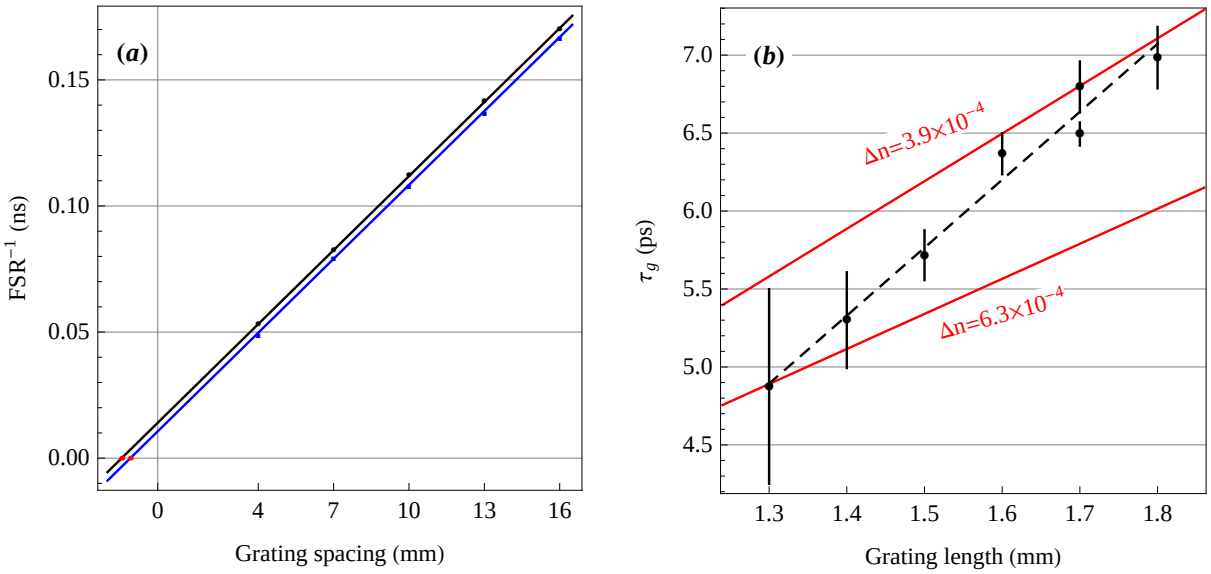

Figure 5.17: **(a)** FSR^{-1} plotted versus length L for five cavities formed between 1.4 mm-long gratings and five using 1.8 mm-long gratings. Determination of group delay by extrapolating to $\mathrm{FSR}^{-1} = 0$, indicated by red dots in the bottom left corner. **(b)** Evolution of group delay as the length of the grating increases. For the penultimate grating, two points are shown: one is derived from the three cavities of Figure 5.18, while the second one uses a different set of data including all five cavities. The dashed line is a linear fit to the data. Red lines: coupled mode theory for two values of index contrast.

5.4.4 Experimental results II: group delay spectroscopy

So far we have only looked at the grating properties at the centre of the stop band. We are now going to have a broader look at the entire grating bandwidth, focusing on the FSR and the group delay. We acquired a new set of spectra on three grating cavities with spacings $L = 4$, 10 and 16 mm between the front faces of the Bragg mirrors (Set E, waveguides 1, 3 and 5). Each mirror is 1.7 mm long, corresponding to approximately 6400 grating planes. The variation of the FSR with frequency for each cavity is shown by the dots in Figures 5.18(a), (b) and (c). As expected the FSR reaches a maximum at the central wavelength of the gratings, where the reflectivity is highest, and therefore the field penetration L_g into the gratings is smallest.

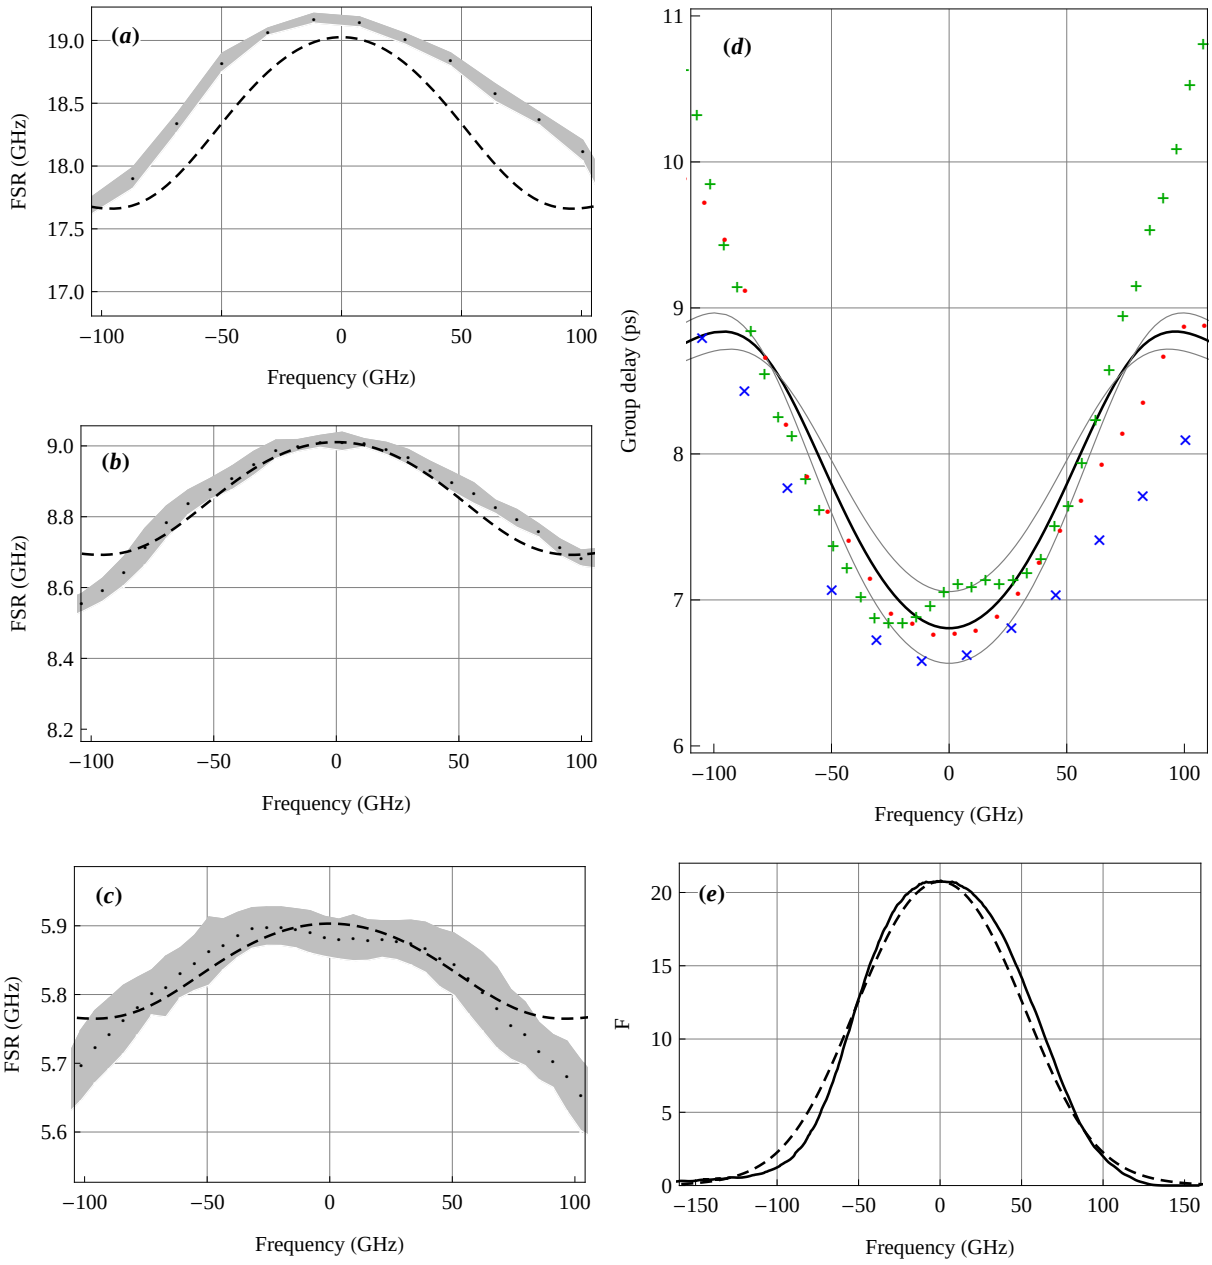

Figure 5.18: (a)-(c): free spectral range versus frequency for cavities of spacing (a) 16 mm, (b) 10 mm, and (c) 4 mm. Points: experimental data with uncertainties shown as shaded area. Dashed line: theoretical FSR derived from coupled mode theory. **(d)** Grating group delay τ_g. Blue crosses, red dots and green pluses are derived from the data in (a), (b) and (c) respectively, while the solid black line is calculated from coupled mode theory fitted to the data with index contrast as the only fit parameter. **(e)**: Measured (solid) and theoretical (dashed line) finesse of the cavity in (b), fitted by adjusting the propagation loss and the index contrast.

We take the frequency at which the FSR peaks to be the Bragg frequency f_B. Averaged over five such cavities we obtain the value $f_B = 3.8422(1) \times 10^{14}\,\mathrm{Hz}$. Together with the known grating spacing $\Lambda_B = 267.3\,\mathrm{nm}$, this yields a mean refractive index $n = 1.4595(1)$ through the relation $n = \frac{c}{2f_B\Lambda_B}$. The group delay $\tau_g = \frac{n}{c}(L_{\mathrm{eff}} - L)$ is then given by

$$\tau_g = \frac{1}{2\mathrm{FSR}} - \frac{nL}{c} \tag{5.14}$$

and is plotted in Figure 5.18(d) for all three cavities. Near the centre of the band, there is good agreement between the results from the three cavities and the theory, represented by the solid line, with $\Delta n = 3.9(5) \times 10^{-4}$ being the only free parameter. The two thin gray lines are the theoretical group delay with $\Delta n = 3.85$ and 3.95×10^{-4} and represent the theoretical error bar; they encompass all the data points near the Bragg frequency. Away from that frequency, however, there are clear systematic differences between the three data sets. These could be due to variation from one grating to another in the apodisation profile. The theory also differs significantly from all three data sets in that it predicts a much faster and more pronounced flattening off of the group delay away from the Bragg frequency. This seems to suggest that the real apodisation profile is not in fact Gaussian, but we do not have enough information to reconstruct what it is.

Finally, we turn to Figure 5.18(e) which shows the measured finesse coefficient F for the 10 mm-long cavity, plotted against frequency (solid line). For a theoretical cavity with grating reflectivity R and single-pass attenuation $\gamma = \exp(-\alpha L_{\mathrm{eff}})$ due to propagation loss, one finds that $F = \frac{4R\gamma}{(1-R\gamma)^2}$ (see derivation in Chapter 3.2). By adjusting γ and changing R through the choice of Δn, we obtain a theoretical curve (dashed) that corresponds reasonably well with the measurements when $\Delta n = 5.5 \times 10^{-4}$ and the absorption coefficient is $\alpha = 0.9\,\mathrm{dB/cm}$. It is difficult to assign error bars here because, once again, there is a systematic difference between theory and experiment in the wings of the band, making this method of measuring Δn inferior to the group delay method described above. The primary purpose of this measurement is to estimate the absorption coefficient α which we consider to be accurate within about 20%.

It is instructive to compare this $\alpha = 0.9(2)\,\mathrm{dB/cm}$ with the $0.235(6)\,\mathrm{dB/cm}$ loss measured in nearly identical waveguides (115) at 1550 nm. This indicates that only a small part or the loss at 1550 nm, 20% or less, can be due to Rayleigh scattering, which scales as λ^{-4}. We believe that the total loss may be significantly reduced in future by improving control over the planar silica layers.

5.4.5 Conclusion

Cavities integrated into optical waveguide chips are of great interest for chemical and biological analysis on a chip and for applications in quantum optics and quantum infor-

mation processing. Visible and very near infrared wavelengths are particularly important for these applications. We have developed such cavities incorporating Bragg reflectors, fabricated by writing in a photosensitive doped silica layer using UV laser light. In the course of characterising them, we have measured transmission spectra from which we derive values for the free spectral range as a function of frequency. We have shown how this provides a simple way to measure the delay of light reflected from the mirrors. Using a coupled mode approach, we have used this delay to determine the index contrast of the gratings. These quantities, delay and index contrast, are of sufficiently general interest that our very simple method may be of use in characterising other devices such as fibre Bragg gratings and delay lines, which are widely used in telecommunications and astronomy. However the reflectivity is too low for applications requiring a high finesse, such as QED experiments.

Chapter 6

Conclusion and prospects

In this thesis we have described our experimental efforts towards the first realisation of a coupled-cavity QED system based on waveguide-coupled Fabry-Pérot microcavities containing atoms — a flexible, tunable, practical platform for quantum simulations and many-body quantum physics. Although no fully working device has been built yet, important steps have been taken. Drawing on existing technology at the University of Southampton, UV-written waveguide chips were fabricated, which were made into optical resonators by depositing multi-layer dielectric mirror on the end facets. A single waveguide resonator was coupled to a free space microcavity formed by one of these mirrors and a spherical mirror etched on a silicon substrate. The latter had been developed for the microcavity QED experiment at Imperial. We observed normal-mode splitting in the reflection spectrum of the coupled waveguide-microcavity system, indicating strong optical coupling between the two. We then coupled two waveguide resonators on the chip, and again we observed normal-mode splitting. We successfully tuned the coupling rate g_{ww} by introducing phase shifts with NiCr heaters deposited on top of the waveguides. This is in itself a new photonic device with potential applications as a tunable delay line in the context of coupled-resonators optical waveguides (139). Finally, in a related but independent project, we demonstrated for the first time Bragg gratings operating at 780 nm in UV-written waveguides and we investigated their properties by combining two gratings into a cavity.

Waveguide loss

It certainly did not escape the reader's perspicacity that in all these experiments the finesse of our waveguide cavities (either formed by Bragg gratings or end-facets DBRs) is rather low — 10 to 40 at most — as a result of large propagation losses. This is why we did not attempt to introduce atoms or other quantum emitters in our system, since we cannot hope to observe quantum effects such as Rabi splitting or Purcell enhancement, much less quantum dynamics, in such a strongly dissipative system. We require the loss

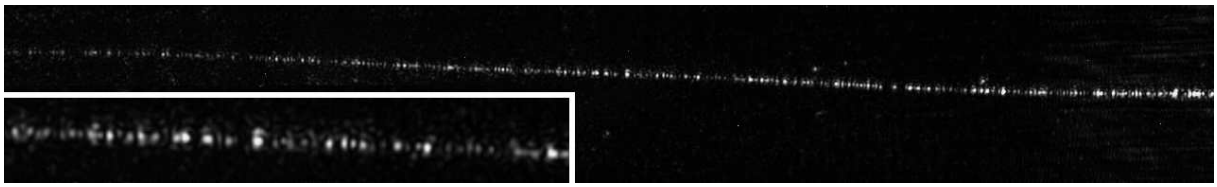

Figure 6.1: Propagation loss in UV-written waveguides. A 30 s CCD exposure with maximum gain reveals that light is being scattered at many localised sites along the waveguides. Inset: detail of the central area

to be, at minimum, not larger than the mirrors largest transmission, say 1% per single waveguide pass. If the waveguide is 1 cm long, this translates into a maximum tolerable loss of 0.01 dB/cm.

Of all the waveguides we used, the femtosecond-written sample from Jena university had the lowest loss, although it is still too large at 0.3 dB/cm. Most of our chips were UV-written, with a loss as high as 1 dB/cm. To investigate its origin we recorded high-gain, long-exposure micrographs of ORC waveguides. A typical example is shown in Figure 6.1 (light is coupled from the right). We see that light is being scattered out of the waveguide at many well-localised spots.

The origin of this high loss is not entirely clear. In any case it is much higher than the intrinsic absorption of silica at 780 nm (optical fibres operating at this wavelength have a propagation loss of 3-4 dB/km, or 0.00003-4 dB/cm). Another common source of loss is the wall roughness, although this is expected to quite small for UV-written waveguides: in one dimension the index profile is defined by the boundary between the core and cladding layer, which is ultra-smooth at small length scales. In the other dimension it is limited by the writing laser, and any imperfection of the laser spot should average out as the sample is translated during the writing process.

In the early developments of the UV-writing technology at the ORC, inhomogeneities in the germanium doping were a major problem. The germanium ions would bunch into islands and wires, so that it was impossible to write a continuous waveguide. Improvements in wafer fabrication solved the problem. However, it could be that residual inhomogeneities contribute to variations of the refractive index along the waveguide, which would induce scattering. If this is indeed the main source of loss at the moment (which is rather difficult to assert accurately), there should be room for improvement, provided one is willing to spend a significant amount of time and money optimising the wafer production. Nonetheless, given that the best "ultra-low loss" UV-written waveguides reported in the literature do not go below 0.2 dB/cm, it is doubtful whether one can reasonably hope to ever improve the loss much below 0.1 dB/cm. The prospects for femtosecond or deposition/etching waveguide are not much better (see Chapter 2.2 for references).

Future directions

It is likely that overcoming the loss problem and building a working waveguide-based coupled-cavity QED platform will require different, emergent technologies. A promising candidate, compatible with the current design, are flat fibres. Currently under development at the ORC (104), they are made from a fibre preform that is flattened during the drawing process. The resulting core layer can be 5 μm thick and a few millimetres wide. If the core had been previously doped with germanium, waveguides can be written exactly as with the silica-on-silicon chips. Fibre-like optical properties can be expected in terms of material purity and homogeneity, resulting in very low loss. On the other hand, the scalability will be more limited: no more than a dozen waveguides or so may fit in the flat area of the fibre.

In another remarkable advance, a propagation loss as low as 0.1 dB/m at 1550 nm was achieved in waveguides consisting of a silicon nitride core with a high aspect ratio (50 nm $\times$ 5 μm) buried in silica (143). Despite the core's aspect ratio, the TE mode is quasi-circular and has a waist of about 3 μm, similar to that of an optical fibre. Even if the propagation loss is significantly larger at 780 nm, it may still be low enough for our purpose.

Other alternatives may require a radical revision of our design. For example, instead of the low contrast waveguides we have used so far (which are very similar to optical fibres), it may be advantageous to use surface (ridge) high-index-contrast waveguides in silicon, silicon nitride or tantalum. These have in general a lower loss thanks to a stronger confinement. Silicon waveguides may integrate bends with a radius of curvature of only tens of microns, whereas we were limited to several millimetres. The downside is that the waveguide mode, being an order of magnitude smaller, is no longer compatible with the microcavities.

One can then follow one of two roads. Either one shrinks the microcavity by decreasing the silicon mirrors' radius of curvature. The group of Jason Smith at Oxford recently reported on microcavity arrays with a radius of curvature as small as 7 μm. Of course in this regime one loses the ability to address the quantum emitters individually from outside the cavity, but at the same time the mode volume is substantially smaller. The other possibility would be to forget about the external cavity altogether, and to integrate resonators on the waveguide itself, using Bragg gratings as mirrors. They could be fabricated by focus ion beam milling. Phase shifters can still be integrated: heating elements if the waveguides are made of silicon or silicon nitride, electro-optics ones if using lithium niobate.

With surface waveguides, dye molecules embedded in a crystalline matrix are prime candidates for the role of quantum emitters. Progress along the lines of Ref. (12) is being made at CCM, and we have also undertaken to integrate dye molecules in the Fabry-Pérot

microcavities.

Appendix A

Optics of Gaussian beams

A.1 Fundamentals

The Gaussian beam is the propagating solution to the wave equation with both the minimum transverse extension and the minimum divergence (144). Its electric field amplitude is:

$$E(x,y,x,t) = \sqrt{\frac{2}{\pi w_0}} \exp\left(-\frac{x^2+y^2}{w^2(z)}\right) \exp\left(-ik\frac{x^2+y^2}{2R(z)} - ikz + i\arctan(z/z_R)\right) \quad \text{(A.1)}$$

and is normalised such that $\iint |E(x,y)|^2 dxdy = 1$. The beam radius $w(z)$ is defined at the $1/e$ half-width of the electric field amplitude, or $1/e^2$ half-width of the intensity $I = |E|^2$, as illustrated in Figure A.1. At the origin it is called the waist w_0, and as a function of propagation distance z we have

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

The Rayleigh length or range $z_R = \pi w_0^2/\lambda$ is the distance at which the radius has increased by a factor $\sqrt{2}$. The wavefront is spherical with a radius of curvature

$$R(z) = z + \frac{z_R^2}{z} \quad \text{(A.3)}$$

which is maximum at the Rayleigh length and flat at the waist and at infinity. The longitudinal (Gouy) phase term introduces an extra π phase shift as the beam goes through the focus: the beam appears to propagate more slowly. It can be neglected for our many wavelengths-long microcavities.

Higher order Gaussian modes can be described by modulating the fundamental mode with Laguerre or Hermite polynomials. For a very good introduction to Gaussian beams see Kogelnik's classic paper (145).

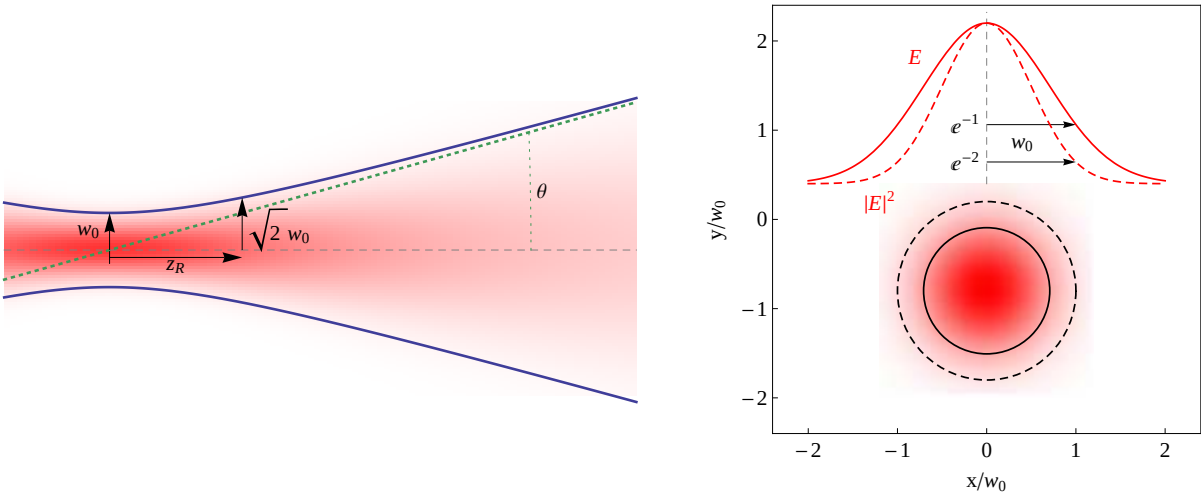

Figure A.1: Longitudinal and transverse mode profile of a Gaussiam beam. The parameters are detailed in the text.

A.2 Mode matching efficiency

In various places throughout this thesis we had to butt-couple Gaussian beams (e.g. fibre-waveguide, waveguide-microcavity). The coupling efficiency η is given by the integral overlap (i.e. the convolution) (146) between the two normalised mode fields functions E and E', which we take here to be cylindrical Gaussians of waists $w_{x,y}$ anf $w'_{x,y}$ whose centres are misaligned by $\Delta x, \Delta y$:

$$\eta = \iint E(x,y)E'^{*}(x-x_0, y-y_0)dxdy \tag{A.4}$$

$$= \iint \frac{2}{\pi\sqrt{w_x w'_x}\sqrt{w_y w'_y}} e^{\left(-\frac{x^2}{w_x^2}-\frac{y^2}{w_y^2}\right)} e^{\left(-\frac{(x-\Delta x)^2}{w'^2_x}+\frac{(y-\Delta y)^2}{w'^2_y}\right)} dxdy \tag{A.5}$$

$$= 2\sqrt{\frac{w_x w_y w'_x w'_y}{(w_x^2+w_x'^2)(w_y^2+w_y'^2)}} \exp\left(-\frac{\Delta x^2}{w_x^2+w_x'^2} - \frac{\Delta y^2}{w_y^2+w_y'^2}\right) \tag{A.6}$$

The two main features of the last expressions are:

1. If the modes are aligned, the intensity mode-matching efficiency (i.e. the power coupled from the first to the second mode) is

$$\eta^2 = 4\frac{w_x w_y w'_x w'_y}{(w_x^2+w_x'^2)(w_y^2+w_y'^2)}. \tag{A.7}$$

2. If the size of the fibre mode E is known, the size of the unknown waveguide E' can be determined by scanning the fibre in front of the waveguide. The detected intensity, as a function of misalignment, will be proportional to η^2, and therefore will be a Gaussian of widths $W_x = \sqrt{w_x^2+w_x'^2)}$ and $W_y = \sqrt{w_y^2+w_y'^2}$. The waveguide

mode size is then $w'_x = \sqrt{W_x^2 - w_x^2}$ and $w'_y = \sqrt{W_y^2 - w_y^2}$. The scanning fibre can be calibrated against itself, we then have $W_{x,y} = \sqrt{2} w_{x,y}$.

A.3 Propagation of Gaussian beam in optical systems

Propagation of Gaussian beams through optical systems can be conveniently described by a lensmaker equation similar to that of geometrical optics:

$$\frac{1}{s_1 + z_{R1}/(s_1 - f)} + \frac{1}{s_2} = \frac{1}{f} \tag{A.8}$$

$s_{1,2}$ and $z_{R1,2}$ are the positions of the object and image waists and the associated Rayleigh ranges; f is the lens focal length. The usual geometric equation is recovered in the "point source" limit $z_R \to 0$. The most surprising feature of Gaussian beam imaging is that the magnification $m = s_2/s_1$ depends on the beam size:

$$m = \frac{s_2}{s_1} = \frac{1}{\sqrt{(1 - s_1/f)^2 + (z_{R1}/f)^2}} \tag{A.9}$$

and $m \to 0$ as $w_0 \to \infty$. This is not so surprising since the magnification's definition breaks down for plane waves.

The second big difference with geometrical optics is that the image is always located at a finite distance from the lens. The maximum image distance is obtained for $s_1 = f + z_{R1}$ and

$$s_{2,\max} = f + \frac{f^2}{2 z_{R1}}. \tag{A.10}$$

The lensmaker equation only relates the image and object waists. A more flexible and powerful is the ABCD matrix formalism, another classic from geometrical optics that can be applied, with some modifications, to Gaussian beams (145). The matrices are the same: for example for free space propagation across a distance s and for a thin lens of focal f we have

$$M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \ , \quad M_{\text{space}} = \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \ , \quad M_{\text{lens}} = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix} \tag{A.11}$$

Complex optical systems are described by multiplying elementary ABCD matrices in the correct order.

We then introduce the beam parameter $q = (1/R - i/z_R)^{-1}$. It is propagated through an arbitrary ABCD matrix as

$$q_2 = \frac{a q_1 + b}{c q_1 + d}. \tag{A.12}$$

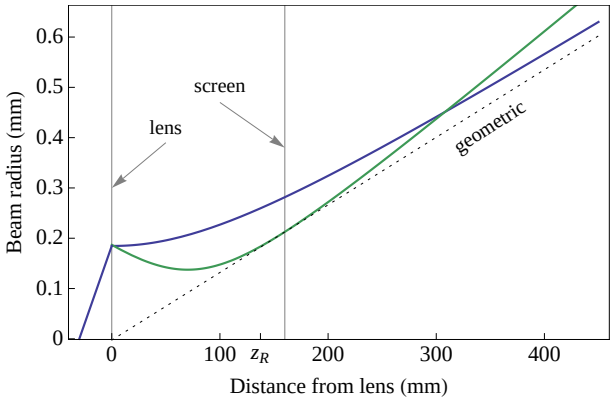

Figure A.2: Microscope imaging of a Gaussian beam. The graph shows the dimensions of the image waist as a function of distance from the lens, $w(z)$. What is the relationship between the geometrical magnification (geometrical image size is represented by the dotted line) and the Gaussian magnification? The screen, located at a fixed distance l from the lens, is indicated by the vertical grey line; it is close to the image Rayleigh distance z_R. This is why an infinity-collimated Gaussian beam (blue curve) will be larger than the geometric image. However the beam will be collimated not at infinity but on the screen (green line); no matter where the screen is, we can always "defocus" the lens so that the beam waist is tangent to the geometrical image at the screen.

We now apply this formalism to the problem of imaging a Gaussian beam with a microscope objective.

Measuring the beam radius

In Chapter 2, we measured the mode radius of a waveguide by imaging its output on a CCD with a high-power ($\times 100$, $f = 1.86\,\mathrm{mm}$) oil immersion objective. The actual magnification, calibrated with a TEM (electron microscope) reference grid, is only 82 because of the non-standard tube length. However, as we've seen previously, the magnification may be different from the geometric one when observing a Gaussian beam. Moreover the problem is compounded by the fact that the image plane is located at a fixed distance from the screen, so that in general we do not record the image waist. We may argue that the rules of geometrical optics should apply since $z_R \approx 20\,\mu\mathrm{m} << f$. On the other hand, the Rayleigh range on the image size is comparable to the screen distance, so Gaussian correction may be required. Let's investigate

In this scenario, the tube length (and therefore the lens-screen distance l) is fixed and we focus the beam by adjusting the object distance s_1 to minimise the image beam radius at the screen. It is tempting to set $s_2 = l$ in the lens maker equation, but a closer examination reveals that for our parameters $s_{2,\max} = 70\,\mathrm{mm} < l$, so that we cannot possibly be imaging the waist. So we have to use the ABCD matrix formalism to find the spot size at an arbitrary distance z_2 from the lens, as a function of object distance s_1:

$$w(z_2) = \frac{w_1}{f z_{R1}} \sqrt{(s_1 z_2 - f(s_1 + z_2))^2 + (f z_{R1} - z_2 z_{R1})^2}. \tag{A.13}$$

Taking the derivative it is then straightforward to show that $w(z_2)$ is minimised when $s_1 = -fz_2/(f-z_2)$ ($s_1 \to f$ when $z_2 \to \infty$, as it should be), and that $w(z_2)_{\min} = w_1(f-z_2)/f$, and the magnification is then simply $m = (f-z_2)/f$. Which is none other than the geometric magnification. Note that nowhere do we have made any approximations: this an exact property of this particular configuration, and is illustrated in Fig. A.2. Thus it is in general correct to use the geometric magnification when imaging a Gaussian beam.

Appendix B

Coupling of a dipole to a nanofibre

In the summer of 2011 I became involved in a project with Jaesuk Hwang and his master student, Nico Verhart from Leiden University, Netherlands. The project consisted in depositing quantum dots on the surface of nanofibres (optical fibres that have been stretched down to a diameter of 1 μm or less), which results in the guided modes extending significantly out of the fibre together with a very small mode volume. This leads to an enhanced coupling between the dipole and the guided mode whereby up to 25% of the dipole's spontaneous emission ends up into a single guided, as has since then been experimentally demonstrated by another group (15).

In addition to the experiment, Nico Verhart was theoretically investigating the dipole-nanofibre coupling by mean of finite-difference time-domain FDTD simulations, whose interpretation was sometimes difficult (especially since we hadn't realised yet that the fibres were multimode), and we decided to look into analytical solutions to the problem. Indeed, according to Fermi's golden rule, the spontaneous emission rate into a given guided mode is given by

$$\Gamma = 2\pi g^2 D(\omega).$$

g is the dipole-field coupling strength, which is related to the mode size and profile, and $D(\omega)$ is the density of modes, which depends on the group velocity. For circular fibre, both quantities can be analytically calculated, which is what I set out to do. First I will briefly discuss the calculation of the nanofibre modes, before showing how to derive useful expressions for the coupling rate to the nanofibre from the golden rule.

B.1 Derivation of coupling rates from Fermi's golden rule

Fermi's golden rule originates from first-order time-dependent perturbation theory (see, for example, Sakurai's textbook (147)). It describes the transition probability between a time-dependent (and often, but nor necessarily, constant) perturbation, of which a

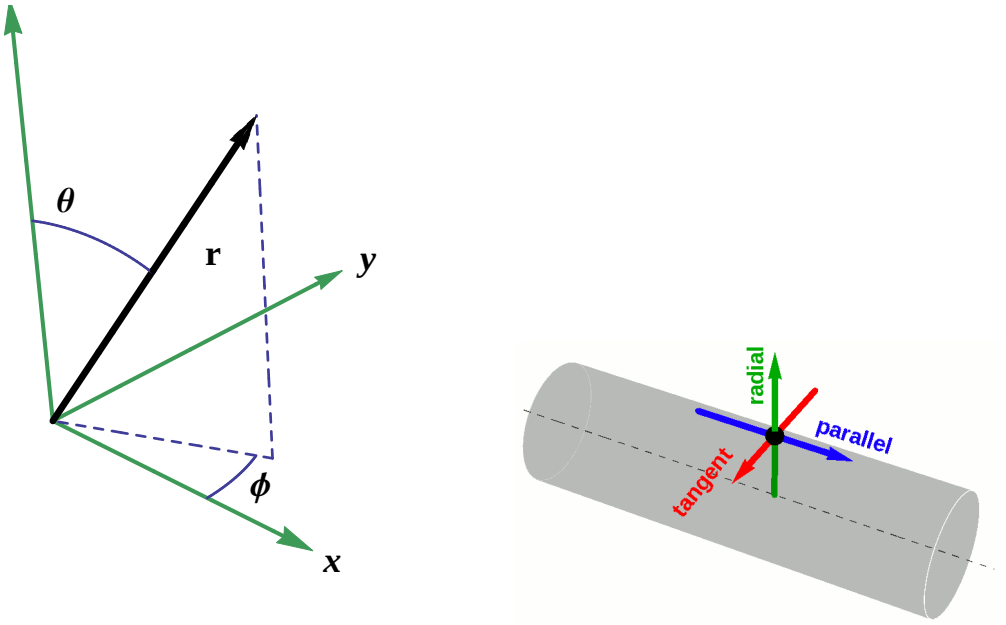

Figure B.1: Left: Spherical coordinate system. Right: the different possible orientations of a dipole near the surface of a nanofibre.

classical electromagnetic field would be a typical and relevant example. Time-dependent essentially means it doesn't have to conserve energy, at least in the limit where the amplitude of the perturbation is not significantly altered during the interaction, and as such it can describe extremely well absorption and emission of radiation between an emitter and the classical field.

What we are interested here, and what the golden rule allows us to do, is to derive an expression for the ratio $\Gamma_{\text{wg}}/\Gamma_{\text{rad}}$ of spontaneous emission into the waveguide vs spontaneous emission in free space. Therefore we have to evaluate this two rates, which involve calculating the density of mode $D(\omega)$ and the coupling strengh g^2 in both cases.

I assume everywhere that the dipole sits in vacuum a short distance away from the fibre, which make the problem somewhat simpler (we don't have to keep track of the refractive index of the medium outside the fibre). Otherwise I follow closely a similar discussion in (12), with some corrections.

B.1.1 Dipole in free space

Density of mode Consider a spherical coordinate system in k-space $\{k,\theta,\phi\}$ as depicted in Figure B.1. The infinitesimal volume element is $dk \times kd\theta \times k\sin\theta d\phi$. Restricting free space to a very large quantisation volume $V = L^3$, each mode occupies a volume $V_0 = (2\pi/L)^3$. To find the density of modes available for a dipole to decay, we have to integrate over a shell $\{k, k+dk\}$ and divide by V_0, taking into account the dipole radiation

pattern $\sin^2\theta$ and the polarisation degeneracy:

$$D(k)dk = 2 \times \left(\frac{L}{2\pi}\right)^3 \times \int_k^{k+dk} k^2 dk \times \int_0^{\pi} \sin^3\theta d\theta \times \int_0^{2\pi} d\phi$$

With $\int_0^{\pi} \sin^3\theta d\theta = 4/3$ and $\int_k^{k+dk} k^2 dk = k^2 dk$, this results in $D(k)dk = \frac{2k^2 dk}{3\pi^2} V$ or, in terms of angular frequency

$$D(\omega)d\omega = \frac{2\omega^2 d\omega}{3\pi^2 c^3} V$$

Coupling strength g Assuming that the dipole is aligned with the field polarisation, g is defined as

$$g^2 = \frac{d^2 E_{\text{vac}}^2}{\hbar^2}$$

where E_{vac} is the vacuum field amplitude. In the same quantisation box as above, it is given by $\hbar\omega = \int \epsilon_0 E_{\text{vac}}^2 d^3V$ or

$$E_{\text{vac}}^2 = \frac{\hbar\omega}{\epsilon_0 V}$$

Note that since we are only interested in the ratio of guided to free space emission, I have dropped annoying factors of 2 that would have cancelled later on anyway. Then

$$g^2 = \frac{d^2\omega}{\hbar\epsilon_0 V}$$

Spontaneous decay rate in free space Using the results from above, we get

$$\Gamma_{\text{rad}} = \frac{4}{3\pi}\frac{d^2}{\hbar\epsilon_0}\frac{\omega^3}{c^3}$$

B.1.2 Dipole on waveguide/nanofibre

Density of modes The k-space is now one-dimensional. With a similar quantisation "volume" L, each mode occupies a volume $2\pi/(Ln_g)$, where n_g is the mode's group index. The density of modes is then simply $D(k)dk = \frac{Ln_g}{2\pi}dk$ for propagation along one direction only (in the 3D case, bi-directional propagation is implicit in the $\int_0^{2\pi} d\phi$ term). In terms of angular frequency this becomes

$$D(\omega) = \frac{Ln_g}{2\pi c} = \frac{L}{2\pi v_g}$$

where $v_g = c/n_g$ is the group velocity. The polarisation here is set by the guided mode, so there is no degeneracy to account for.

Waveguide-dipole coupling strength The mode is described by an electric field amplitude $\mathbf{E}(\mathbf{r})$, and the waveguide by its relative permittivity $\epsilon_r(\mathbf{r}) = n^2(\mathbf{r})$. The coupling

strength is

$$g^2 = d^2 E_{\text{vac}}^2(\mathbf{r}_0)/\hbar^2$$

where $\mathbf{r}_0$ is the position of the molecule. The vacuum field is such that

$$\hbar\omega = L \int \epsilon_0 \epsilon_r(\mathbf{r}) E_{\text{vac}}^2(\mathbf{r}) d^2\mathbf{r}$$

where L is the same quantisation length as before. It is then convenient to define an effective mode area

$$A_{\text{eff}} = \frac{\iint \epsilon_r(\mathbf{r}) E^2(\mathbf{r}) d^2\mathbf{r}}{E^2(\mathbf{r}_0)}$$

which can be easily calculated and recast in term of the vacuum field as

$$A_{\text{eff}} = \frac{\hbar\omega/(\epsilon_0 L)}{E_{\text{vac}}^2(\mathbf{r}_0)}.$$

Substituting $E_{\text{vac}}^2(\mathbf{r}_0)$ in g^2 then gives simply

$$g^2 = \frac{d^2\omega}{\hbar\epsilon_0 L A_{\text{eff}}}$$

Spontaneous emission rate into waveguide Combining the two previous results,

$$\Gamma_{\text{wg}} = \frac{d^2\omega}{\hbar\epsilon_0 A_{\text{eff}} v_g}$$

B.1.3 Guided to free space ratios

The ratio of spontaneous emission into the waveguide (one direction only) to spontaneous emission in free space is

$$\frac{\Gamma_{\text{wg}}}{\Gamma_{\text{rad}}} = \frac{3\pi}{4A_{\text{eff}} v_g} \frac{c^3}{\omega^2} = \frac{3c\lambda^2}{16\pi A_{\text{eff}} v_g}.$$

Note that this is different from the fraction of spontaneous emission coupled to the fibre relative to the total spontaneous emission in presence of the fibre, which is much more difficult to calculate as it involves a large number of modes of different nature.

B.2 Modes of the nanofibre

We have already gone over the procedure to calculate the mode of an optical fibre in Chapter 2. Because of the high index contrast ($n_{\text{co}} - n_{\text{cl}} = n_{\text{SiO}_2} - 1 = 0.45$), the weak guiding approximation does not apply: we have to use hybrid modes. The dispersion relation of the first few modes as a function of the nanofibre radius is plotted in Figure B.2(a) (the wavelength is set to 605 nm), whereas in Fig. B.2(b) we plotted for comparison the dispersion relation of the initial, non-tapered fibre (SMF28, index contrast 0.0067, core

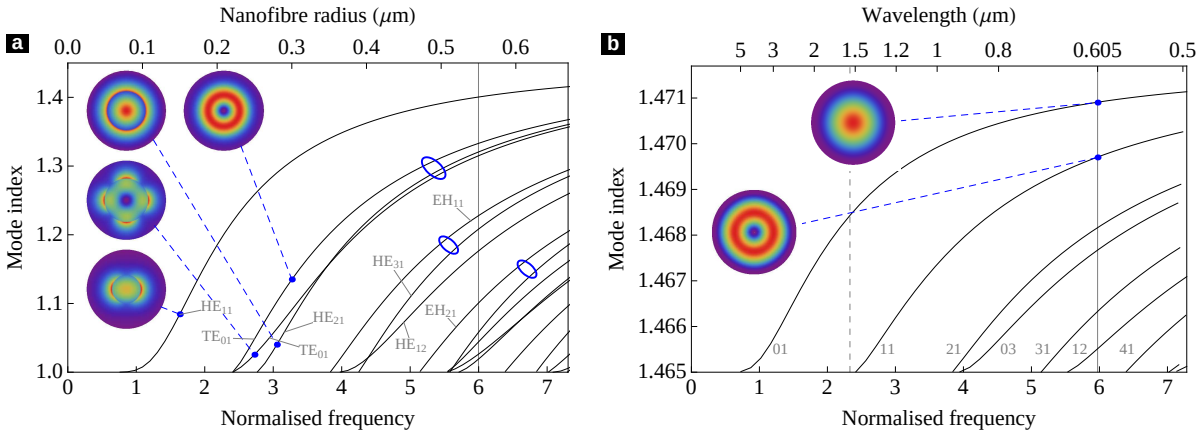

Figure B.2: Left: Dispersion diagram and selected hybrid mode profiles of a nanofibre. Notice how the field can extend significantly outside the core. Right: same diagram but for a standard fibre in the weak guiding (LP modes) regime. A standard single-mode fibre for 1550 nm (SMF28) lies close to $v = 6$ (vertical line).

diameter 8.2 μm) as a function of the wavelength. The transverse mode profiles of a few selected modes are also shown. We can see that the nanofibre mode extend quite significantly outside the core; in particular looking at the EH_{11} mode, we can see that the field peaks at two diametrically opposed points just outside the core. The effective mode area at these maxima is only $0.62\,\lambda^2$ (at $\lambda = 605$ nm) for a radially-oriented dipole, so we can therefore expect that an emitter located at one of this point would experience a strong coupling to the waveguide.

We also need the group index/velocity. This can be calculated from the modal index as $n_g = n_{\mathrm{eff}} - \lambda_0 dn/d\lambda_0$ with λ_0 the vacuum wavelength. Alternatively, it can be calculated from the modal fields via the elegant theorem (74)

$$\mathbf{v}_g = \frac{\int \mathbf{S} d^3\mathbf{r}}{U_E + U_H} \tag{B.1}$$

$$\mathbf{S} = \frac{1}{2}\mathrm{Re}[\mathbf{E}^* \times \mathbf{H}] \tag{B.2}$$

$$U_E = \frac{\varepsilon_0}{4}\int \varepsilon(\mathbf{r})|\mathbf{E}(\mathbf{r})|^2 d^3\mathbf{r} \tag{B.3}$$

$$U_H = \frac{\mu_0}{4}\int |\mathbf{H}(\mathbf{r})|^2 d^3\mathbf{r} \tag{B.4}$$

where $\mathbf{S}$, U_E and U_H are respectively the Poynting vector and the energy stored in the electric and magnetic fields (of course, $U_E = U_H$). In other words, the velocity of energy propagation is the ratio of energy flux to energy density.

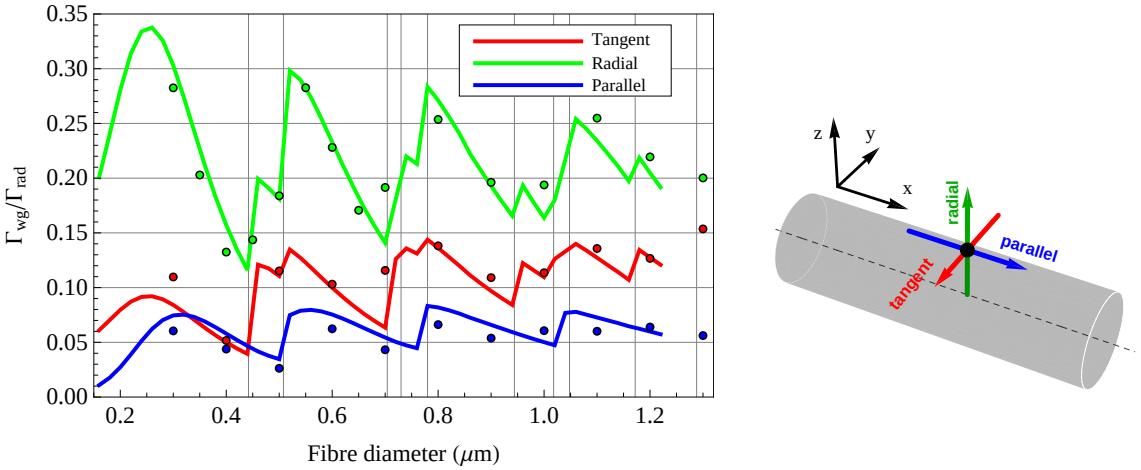

Figure B.3: Spontaneous decay rate of an atom (relative to free space emission γ_0) sitting 10 nm above a nanofibre as a function of the fibre diameter. The inset describes the three possible orientations of the dipole. The coupling tends to decrease as the fibre gets larger, but jumps upward whenever new decay channels become available in the form of additional guided modes (grey vertical lines are the cut-off frequencies).

B.3 Discussion

We know have all the tools needed to compute the fraction of emission coupled to the fibre, $\Gamma_{\rm wg}/\Gamma_{\rm rad}$, as a function of the nanofibre diameter. If the fibre is large enough to support several modes, we only have to add the contribution of the different decay channels: since all guided modes are orthogonal, they can be treated independently.

The final result are plotted in Figure B.3. It shows the $\Gamma_{\rm wg}/\Gamma_{\rm rad}$ ratio as a function of the nanofibre diameter, together with the FDTD results. The two agree remarkably well, considering how difficult it can be to extract the guided vs total power from FDTD simulations (largely because the computational cell is limited to a few tens of microns along the fibre due to memory constrains). We observe an oscillatory structure with periods of exponential decays as the modes confinement increases with the fibre diameter, with sudden jumps coinciding with the appearance of new guided modes offering additional decay channels to the emitter. It is quite remarkable that once these oscillations are abstracted, the coupling ratio is essentially constant.

It can be seen that the coupling is significantly larger for radially-oriented dipoles. This is because the electric field perpendicular to a dielectric interface is not continuous (unlike the tangential components of the field), and the magnitude of the discontinuity is given by the refractive index difference (in this case, $n_{\rm co} - n_{\rm cl} = 0.45$). The radial component of the field, and the coupling of radial dipoles, is thus enhanced by about 50% compared to the tangential components.

The maximum coupling occurs for a nanofibre diameter of about 300 nm, where it is higher than 30%.

In the multimode case ($V > 2.405$), The collection efficiency is not determined solely

by the nanofibre: higher-order modes also have to survive in the non-tapered section. For a standard SMF28 fibre designed for operation at 1.5 μm, the normalised frequency at 605 nm is $V = 5.98$, very close to that in the tapered region of the fibres used in our experiments ($2a = 1.1\,\mu$m, $V \simeq 6.0$).

Appendix C

Conversion between different loss parameters

In the course of this thesis we have not been very consistent in describing propagation loss in waveguides, preferring to use the most "natural" way in different settings. In this appendix we summarise the different ways of introducing the propagation loss and we give conversion formulae between them.

Decibels per unit length

Consider a waveguide of length l. The optical power at the input and output are P_{in} and P_{out}. The loss in dB per unit length (usually per centimenter) is then defined as

$$L_{dB} = \frac{10}{l} \log_{10} \left(\frac{P_{out}}{P_{in}} \right). \tag{C.1}$$

Cavity imaginary round trip phase Γ

This has been described in section 3.2. To the cavity round trip phase $\phi = 2kL_{\text{cav}}$ we add an imaginary component, so it becomes $\phi - i\Gamma$. L_{cav} is the cavity length. We showed in Chapter 3.2.1 that if the power is P_{in} at the cavity input, it becomes $P_{out} = e^{-\Gamma} P_{in}$ after a single pass. Therefore the dB/cm loss is

$$L_{\text{dB}} = 10 \log_{10} \left(e^{-\Gamma} \right) \frac{1}{L_{\text{cav}}} = -\frac{10}{\ln(10)} \frac{\Gamma}{L_{\text{cav}}}. \tag{C.2}$$

Imaginary refractive index q

Whenever the refractive index appears explicitly, we can introduce an imaginary part q: $\tilde{n} = n + iq$. Consider then a plane wave of vacuum wave vector $k_0 = 2\pi/\lambda$. It propagates

along z as

$$E(z) \propto e^{ik_0(n+iq)z+i\omega t}$$

Therefore its intensity decreases as $e^{-2k_0ql} = e^{-\equiv\Gamma}$ over a distance l, so that the loss in dB/cm is

$$L_{dB} = 10\log_{10}\left(e^{-2kql}\right)\frac{1}{l} = -\frac{20}{\ln(10)}k_0q \tag{C.3}$$

and we have also

$$\Gamma = 2qL_{\text{cav}}k_0 \tag{C.4}$$

Exponential decay rate q

Finally, in section 5.3.1 we introduced a loss parameter in the coupled-mode equations. In the absence of cross- and self-coupling, these become

$$\frac{dA(z)}{dz} = -\gamma A(z) \Rightarrow A(z) = A(0)e^{-\gamma z} \tag{C.5}$$

and therefore, after a length l we have

$$e^{-\gamma l} = e^{-\Gamma} \Rightarrow \gamma l = \Gamma \tag{C.6}$$

Linear dB loss

For small loss (ie $P_{out}/P_{in} > 90\%$), we can write the dB loss as $10\log_{10}(1 - P_{\text{lost}})$ where $P_{\text{lost}} = P_{out}/P_{in}$. Eq. (C.1) can be linearised as

$$L_{\text{dB}} = \frac{10}{\ln 10}P_{\text{lost}} = 4.34P_{\text{lost}} \tag{C.7}$$

Summary

Table C.1 summarises the relationships between all these parameters.

	Γ	q	γ	dB/cm
Γ		$2L_{\mathrm{cav}}k_0q$	γL_{cav}	$-\frac{\ln(10)}{10}L_{\mathrm{cav}}L_{\mathrm{dB}}$
q	$\frac{\Gamma}{2L_{\mathrm{cav}}k_0}$		$\frac{\gamma}{2k_0}$	$\frac{\ln(10)}{20}k_0L_{\mathrm{dB}}$
γ	$\frac{\Gamma}{L_{\mathrm{cav}}}$	$2k_0q$		$-\frac{\ln(10)}{10}L_{\mathrm{dB}}$
dB/cm	$-\frac{10}{\ln(10)}\frac{\Gamma}{L_{\mathrm{cav}}}$	$-\frac{20}{\ln(10)}k_0q$	$-\frac{10}{\ln(10)}\gamma$	

Table C.1: Loss conversion table.

Bibliography

[1] X. Zhou, I. Dotsenko, B. Peaudecerf, T. Rybarczyk, C. Sayrin, S. Gleyzes, J. Raimond, M. Brune, and S. Haroche, "Field Locked to a Fock State by Quantum Feedback with Single Photon Corrections," *Physical Review Letters*, vol. 108, p. 243602, June 2012.

[2] G. Zumofen, N. Mojarad, V. Sandoghdar, and M. Agio, "Perfect Reflection of Light by an Oscillating Dipole," *Physical Review Letters*, vol. 101, pp. 1–4, Oct. 2008.

[3] S. van Enk, "Atoms, dipole waves, and strongly focused light beams," *Physical Review A*, vol. 69, p. 043813, Apr. 2004.

[4] R. Maiwald, D. Leibfried, J. Britton, J. C. Bergquist, G. Leuchs, and D. J. Wineland, "Stylus ion trap for enhanced access and sensing," *Nature Physics*, vol. 5, pp. 551–554, June 2009.

[5] J. Hwang, M. Pototschnig, R. Lettow, G. Zumofen, A. Renn, S. Götzinger, and V. Sandoghdar, "A single-molecule optical transistor," *Nature*, vol. 460, pp. 76–80, July 2009.

[6] M. Pototschnig, Y. Chassagneux, J. Hwang, G. Zumofen, A. Renn, and V. Sandoghdar, "Controlling the Phase of a Light Beam with a Single Molecule," *Physical Review Letters*, vol. 107, pp. 1–5, Aug. 2011.

[7] K. G. Lee, X.-W. Chen, H. Eghlidi, P. Kukura, R. Lettow, A. Renn, V. Sandoghdar, and S. Götzinger, "A planar dielectric antenna for directional single-photon emission and near-unity collection efficiency," *Nature Photonics*, vol. 5, pp. 166–169, Jan. 2011.

[8] R. Lettow, Y. L. A. Rezus, A. Renn, G. Zumofen, E. Ikonen, S. Götzinger, and V. Sandoghdar, "Quantum Interference of Tunably Indistinguishable Photons from Remote Organic Molecules," *Physical Review Letters*, vol. 104, pp. 26–29, Mar. 2010.

[9] Y. L. A. Rezus, S. Walt, R. Lettow, A. Renn, G. Zumofen, S. Götzinger, and V. Sandoghdar, "Single-Photon Spectroscopy of a Single Molecule," *Physical Review Letters*, vol. 108, pp. 1–5, Feb. 2012.

[10] M. Steiner, F. Schleifenbaum, C. Stupperich, A. Virgilio Failla, A. Hartschuh, and A. J. Meixner, "Microcavity-controlled single-molecule fluorescence.," *ChemPhysChem*, vol. 6, pp. 2190–6, Oct. 2005.

[11] P. Londero, J. Levy, A. Slepkov, A. Bhagwat, K. Saha, V. Venkataraman, M. Lipson, and A. L. Gaeta, "Chip-Based Optical Interactions with Rubidium Vapor," in *Proceedings of CLEO 2010*, pp. 4–5, 2010.

[12] J. Hwang and E. a. Hinds, "Dye molecules as single-photon sources and large optical nonlinearities on a chip," *New Journal of Physics*, vol. 13, p. 085009, Aug. 2011.

[13] M. J. Morrissey, K. Deasy, Y. Wu, S. Chakrabarti, and S. Nic Chormaic, "Tapered optical fibers as tools for probing magneto-optical trap characteristics.," *Review of scientific instruments*, vol. 80, p. 053102, May 2009.

[14] S. Dawkins, R. Mitsch, D. Reitz, E. Vetsch, and A. Rauschenbeutel, "Dispersive Optical Interface Based on Nanofiber-Trapped Atoms," *Physical Review Letters*, vol. 107, pp. 1–4, Dec. 2011.

[15] R. Yalla, F. Le Kien, M. Morinaga, and K. Hakuta, "Efficient Channeling of Fluorescence Photons from Single Quantum Dots into Guided Modes of Optical Nanofiber," *Physical Review Letters*, vol. 109, p. 5, Aug. 2012, 1204.6153.

[16] F. Le Kien and K. Hakuta, "Triggered generation of single guided photons from a single atom in a nanofiber cavity," *Physical Review A*, vol. 83, pp. 1–11, Apr. 2011.
[17] C. Wuttke, M. Becker, S. Brückner, M. Rothhardt, and A. Rauschenbeutel, "Nanofiber Fabry-Perot microresonator for nonlinear optics and cavity quantum electrodynamics.," *Optics letters*, vol. 37, pp. 1949–51, June 2012.
[18] A. Faraon, P. E. Barclay, C. Santori, K.-M. C. Fu, and R. G. Beausoleil, "Resonant enhancement of the zero-phonon emission from a colour centre in a diamond cavity," *Nature Photonics*, vol. 5, pp. 301–305, Apr. 2011.
[19] T. Peyronel, O. Firstenberg, Q.-Y. Liang, S. Hofferberth, A. V. Gorshkov, T. Pohl, M. D. Lukin, and V. Vuletić, "Quantum nonlinear optics with single photons enabled by strongly interacting atoms.," *Nature*, vol. 488, pp. 57–60, Aug. 2012.
[20] E. Jaynes and F. Cummings, "Comparison of quantum and semiclassical radiation theories with application to the beam maser," *Proceedings of the IEEE*, vol. 51, no. 1, pp. 89–109, 1963.
[21] R. Dicke, "Coherence in Spontaneous Radiation Processes," *Physical Review*, vol. 93, pp. 99–110, Jan. 1954.
[22] P. Milonni and P. Knight, "Spontaneous emission between mirrors," *Optics Communications*, vol. 9, pp. 119–122, Oct. 1973.
[23] A. Kuhn and G. Rempe, "Optical Cavity QED: Fundamentals and Application as a Single-Photon Light Source," in *Experimental Quantum Computation and Information* (F. Demartini and C. Monroe, eds.), pp. 37–66, IOS-Press, 2002.
[24] K. M. Birnbaum, A. Boca, R. Miller, A. D. Boozer, T. E. Northup, and H. J. Kimble, "Photon blockade in an optical cavity with one trapped atom.," *Nature*, vol. 436, pp. 87–90, July 2005.
[25] M. Fleischhauer, A. Imamoglu, and J. Marangos, "Electromagnetically induced transparency: Optics in coherent media," *Reviews of Modern Physics*, vol. 77, no. 2, p. 633, 2005.
[26] P. B. R. Nisbet-Jones, J. Dilley, D. Ljunggren, and A. Kuhn, "Highly efficient source for indistinguishable single photons of controlled shape," *New Journal of Physics*, vol. 13, p. 103036, Oct. 2011.
[27] E. Purcell, H. Torrey, and R. Pound, "Resonance Absorption by Nuclear Magnetic Moments in a Solid," *Physical Review*, vol. 69, pp. 37–38, Jan. 1946.
[28] F. Bloch, "Nuclear Induction," *Physical Review*, vol. 70, pp. 460–474, Oct. 1946.
[29] M. Gerstein, "Purcell's role in the discovery of nuclear magnetic resonance: Contingency versus inevitability," *American Journal of Physics*, vol. 62, no. 7, p. 595, 1994.
[30] E. Purcell, "Spontaneous emission probabilities at radio frequencies," *Physical Review*, vol. 69, p. 681, 1946.
[31] "Proceedings of the American Physical Society," in *Physical Review*, vol. 69, pp. 674–674, June 1946.
[32] D. Kleppner, "Inhibited Spontaneous Emission," *Physical Review Letters*, vol. 47, pp. 233–236, July 1981.
[33] K. Drexhage, "Interaction of Light with Monomolecular Dye Layers," *Progress in Optics*, vol. 12, no. IV, p. 165, 1974.
[34] K. Drexhage, "Monomolecular Layers and Light," *Scientific American*, vol. 222, pp. 108–119, Mar. 1970.
[35] K. Drexhage, "Influence of a dielectric interface on fluorescence decay time," *Journal of Luminescence*, vol. 1-2, pp. 693–701, Jan. 1970.
[36] P. Goy, J. Raimond, M. Gross, and S. Haroche, "Observation of cavity-enhanced single-atom spontaneous emission," *Physical Review letters*, vol. 50, no. 24, pp. 1903–1906, 1983.
[37] R. Hulet, E. Hilfer, and D. Kleppner, "Inhibited Spontaneous Emission by a Rydberg Atom," *Physical Review Letters*, vol. 55, pp. 2137–2140, Nov. 1985.

[38] G. Gabrielse and H. Dehmelt, "Observation of inhibited spontaneous emission," *Physical Review Letters*, vol. 55, pp. 67–70, July 1985.
[39] D. Meschede, H. Walther, and G. Müller, "One-Atom Maser," *Physical Review Letters*, vol. 54, pp. 551–554, Feb. 1985.
[40] G. Rempe, H. Walther, and N. Klein, "Observation of quantum collapse and revival in a one-atom maser," *Physical Review Letters*, vol. 58, pp. 353–356, Jan. 1987.
[41] W. Jhe, A. Anderson, E. A. Hinds, D. Meschede, L. Moi, and S. Haroche, "Suppression of spontaneous decay at optical frequencies: Test of vacuum-field anisotropy in confined space," *Physical review letters*, vol. 58, no. 7, pp. 666–669, 1987.
[42] R. Thompson, G. Rempe, and H. J. Kimble, "Observation of normal-mode splitting for an atom in an optical cavity," *Physical Review Letters*, vol. 68, pp. 1132–1135, Feb. 1992.
[43] Q. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble, "Measurement of Conditional Phase Shifts for Quantum Logic," *Physical Review Letters*, vol. 75, pp. 4710–4713, Dec. 1995.
[44] C. J. Hood, T. W. Lynn, A. C. Doherty, A. S. Parkins, and H. J. Kimble, "The Atom-Cavity Microscope: Single Atoms Bound in Orbit by Single Photons," *Science*, vol. 287, pp. 1447–1453, Feb. 2000.
[45] J. McKeever, A. Boca, a. D. Boozer, J. R. Buck, and H. J. Kimble, "Experimental realization of a one-atom laser in the regime of strong coupling.," *Nature*, vol. 425, pp. 268–71, Sept. 2003.
[46] J. McKeever, J. R. Buck, A. D. Boozer, a. Kuzmich, H.-C. Nägerl, D. M. Stamper-Kurn, and H. J. Kimble, "State-Insensitive Cooling and Trapping of Single Atoms in an Optical Cavity," *Physical Review Letters*, vol. 90, pp. 2–5, Apr. 2003.
[47] M. Karski, L. Förster, J.-M. Choi, A. Steffen, W. Alt, D. Meschede, and A. Widera, "Quantum walk in position space with single optically trapped atoms.," *Science*, vol. 325, pp. 174–7, July 2009.
[48] S. Brakhane, W. Alt, T. Kampschulte, M. Martinez-Dorantes, R. Reimann, S. Yoon, A. Widera, and D. Meschede, "Bayesian Feedback Control of a Two-Atom Spin-State in an Atom-Cavity System," *Physical Review Letters*, vol. 109, pp. 1–5, Oct. 2012.
[49] M. Pelton, C. Santori, W. von der Linden, B. Zhang, G. S. Solomon, J. Plant, and Y. Yamamoto, "Efficient Source of Single Photons: A Single Quantum Dot in a Micropost Microcavity," *Physical Review Letters*, vol. 89, pp. 1–4, Nov. 2002.
[50] M. Hartmann, F. Brandão, and M. Plenio, "Strongly interacting polaritons in coupled arrays of cavities," *Nature Physics*, vol. 2, pp. 849–855, Nov. 2006.
[51] M. Fisher, P. Weichman, G. Grinstein, and D. Fisher, "Boson localization and the superfluid-insulator transition," *Physical Review B*, vol. 40, no. 1, p. 546, 1989.
[52] M. Greiner, O. Mandel, T. Esslinger, T. Hänsch, and I. Bloch, "Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms," *Nature*, vol. 415, no. 6867, pp. 39–44, 2002.
[53] D. Rossini and R. Fazio, "Mott-Insulating and Glassy Phases of Polaritons in 1D Arrays of Coupled Cavities," *Physical Review Letters*, vol. 99, pp. 2–5, Oct. 2007.
[54] F. Brandão, M. Hartmann, and M. Plenio, "Light-shift-induced photonic nonlinearities," *New Journal of Physics*, vol. 10, p. 043010, Apr. 2008.
[55] A. Imamoglu, H. Schmidt, G. Woods, and M. Deutsch, "Strongly Interacting Photons in a Nonlinear Cavity," *Physical Review Letters*, vol. 79, pp. 1467–1470, Aug. 1997.
[56] P. Grangier, D. Walls, and K. Gheri, "Comment on "Strongly Interacting Photons in a Nonlinear Cavity"," *Physical Review Letters*, vol. 81, pp. 2833–2833, Sept. 1998.
[57] R. Tan, G.-x. Li, and Z. Ficek, "Cavity-induced giant Kerr nonlinearities in a driven V -type atom," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 42, p. 055507, Mar. 2009.
[58] A. Greentree, J. Vaccaro, S. Echaniz, A. Durrant, and J. Marangos, "Prospects for photon

blockade in four-level systems in the N configuration with more than one atom," *Journal of Optics B: Quantum and Semiclassical Optics*, vol. 2, p. 252, 2000.
[59] A. Greentree, C. Tahan, J. Cole, and L. C. L. Hollenberg, "Quantum phase transitions of light," *Nature Physics*, vol. 2, pp. 856–861, Nov. 2006.
[60] M. Hartmann and M. Plenio, "Strong Photon Nonlinearities and Photonic Mott Insulators," *Physical Review Letters*, vol. 99, Sept. 2007.
[61] M. Hartmann, F. Brandão, and M. Plenio, "Effective Spin Systems in Coupled Microcavities," *Physical Review Letters*, vol. 99, pp. 19–22, Oct. 2007.
[62] D. Angelakis, M. Santos, and S. Bose, "Photon-blockade-induced Mott transitions and XY spin models in coupled cavity arrays," *Physical Review A*, vol. 76, pp. 3–6, Sept. 2007.
[63] S. Giampaolo and F. Illuminati, "Long-distance entanglement and quantum teleportation in coupled-cavity arrays," *Physical Review A*, vol. 80, pp. 4–7, Nov. 2009.
[64] M. Hartmann, F. Brandão, and M. Plenio, "Quantum many-body phenomena in coupled cavity arrays," *Laser & Photonics Review*, vol. 2, pp. 527–556, Dec. 2008.
[65] M. Knap, E. Arrigoni, W. von der Linden, and J. Cole, "Emission characteristics of laser-driven dissipative coupled-cavity systems," *Physical Review A*, vol. 83, pp. 1–11, Feb. 2011.
[66] M. Hartmann, "Polariton Crystallization in Driven Arrays of Lossy Nonlinear Resonators," *Physical Review Letters*, vol. 104, p. 113601, Mar. 2010.
[67] I. Carusotto, D. Gerace, H. Tureci, S. De Liberato, C. Ciuti, and A. Imamoglu, "Fermionized Photons in an Array of Driven Dissipative Nonlinear Cavities," *Physical Review Letters*, vol. 103, p. 033601, July 2009.
[68] S. Ferretti, L. Andreani, H. Türeci, and D. Gerace, "Photon correlations in a two-site nonlinear cavity system under coherent drive and dissipation," *Physical Review A*, vol. 82, p. 013841, July 2010.
[69] R. Umucalılar and I. Carusotto, "Fractional Quantum Hall States of Photons in an Array of Dissipative Coupled Cavities," *Physical Review Letters*, vol. 108, p. 206809, May 2012.
[70] M. Trupke, E. A. Hinds, S. Eriksson, E. A. Curtis, Z. Moktadir, E. Kukharenka, and M. Kraft, "Microfabricated high-finesse optical cavity with open access and small volume," *Applied Physics Letters*, vol. 87, no. 21, p. 211106, 2005.
[71] M. Trupke, J. Goldwin, B. Darquié, G. Dutier, S. Eriksson, J. Ashmore, and E. A. Hinds, "Atom Detection and Photon Production in a Scalable, Open, Optical Microcavity," *Physical Review Letters*, vol. 99, p. 063601, Aug. 2007.
[72] J. Goldwin, M. Trupke, J. Kenner, a. Ratnapala, and E. a. Hinds, "Fast cavity-enhanced atom detection with low noise and high fidelity.," *Nature communications*, vol. 2, p. 418, Jan. 2011.
[73] M. Kohnen, M. Succo, P. G. Petrov, R. A. Nyman, M. Trupke, and E. A. Hinds, "An array of integrated atom–photon junctions," *Nature Photonics*, vol. 5, pp. 35–38, Nov. 2010.
[74] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic crystals*, vol. 102. Princeton University Press, 2008.
[75] M. Notomi, E. Kuramochi, and T. Tanabe, "Large-scale arrays of ultrahigh-Q coupled nanocavities," *Nature Photonics*, vol. 2, pp. 741–747, Nov. 2008.
[76] Y. Sato, Y. Tanaka, J. Upham, Y. Takahashi, T. Asano, and S. Noda, "Strong coupling between distant photonic nanocavities and its dynamic control," *Nature Photonics*, vol. 6, pp. 56–61, Dec. 2011.
[77] D. Englund, A. Majumdar, M. Bajcsy, A. Faraon, P. Petroff, and J. Vucković, "Ultrafast Photon-Photon Interaction in a Strongly Coupled Quantum Dot-Cavity System," *Physical Review Letters*, vol. 108, Mar. 2012.
[78] R. Bose, T. Cai, G. S. Solomon, and E. Waks, "All-optical tuning of a quantum dot in a coupled cavity system," *Applied Physics Letters*, vol. 100, no. 23, p. 231107, 2012.
[79] D. K. Armani, T. J. Kippenberg, S. M. Spillane, and K. J. Vahala, "Ultra-high-Q toroid

microcavity on a chip.," *Nature*, vol. 421, pp. 925–8, Feb. 2003.
[80] A. Mazzei, S. Götzinger, L. de S. Menezes, G. Zumofen, O. Benson, and V. Sandoghdar, "Controlled Coupling of Counterpropagating Whispering-Gallery Modes by a Single Rayleigh Scatterer: A Classical Problem in a Quantum Optical Light," *Physical Review Letters*, vol. 99, Oct. 2007.
[81] P. E. Barclay, K. Srinivasan, O. Painter, B. Lev, and H. Mabuchi, "Integration of fiber-coupled high-Q SiNx microdisks with atom chips," *Applied Physics Letters*, vol. 89, no. 13, p. 131108, 2006.
[82] D. Reitz and A. Rauschenbeutel, "Nanofiber-based double-helix dipole trap for cold neutral atoms," *Optics Communications*, vol. 285, pp. 4705–4708, Oct. 2012.
[83] R. Amsüss, C. Koller, T. Nöbauer, S. Putz, S. Rotter, K. Sandner, S. Schneider, M. Schramböck, G. Steinhauser, H. Ritsch, J. Schmiedmayer, and J. Majer, "Cavity QED with Magnetically Coupled Collective Spin States," *Physical Review Letters*, vol. 107, pp. 1–5, Aug. 2011.
[84] J. Majer, J. M. Chow, J. M. Gambetta, J. Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Coupling superconducting qubits via a cavity bus.," *Nature*, vol. 449, pp. 443–7, Sept. 2007.
[85] "Hybrid Information Processing website." http://www.dmi.unisa.it/HIP/index.html.
[86] M. Trupke, J. Metz, A. Beige, and E. A. Hinds, "Towards quantum computing with single atoms and optical cavities on atom chips," *Journal of Modern Optics*, vol. 54, pp. 1639–1655, July 2007.
[87] G. W. Biedermann, F. M. Benito, K. M. Fortier, D. L. Stick, T. K. Loyd, P. D. D. Schwindt, C. Y. Nakakura, R. L. Jarecki, and M. G. Blain, "Ultrasmooth microfabricated mirrors for quantum information," *Applied Physics Letters*, vol. 97, no. 18, p. 181110, 2010.
[88] "Oz Optics." http://www.ozoptics.com.
[89] G. Lepert, M. Trupke, M. Hartmann, M. Plenio, and E. a. Hinds, "Arrays of waveguide-coupled optical cavities that interact strongly with atoms," *New Journal of Physics*, vol. 13, p. 113002, Nov. 2011.
[90] E. Snitzer, "Cylindrical Dielectric Waveguide Modes," *Journal of the Optical Society of America*, vol. 51, p. 491, May 1961.
[91] D. Gloge, "Weakly Guiding Fibers," *Applied Optics*, vol. 10, p. 2252, Oct. 1971.
[92] J. E. Goell, "A circular-Harmonic computer analysis of rectangular dielectric waveguides," *Bell Syst. Tech J.*, vol. 48, no. 7, p. 2133, 1969.
[93] M. I. Nathan, W. P. Dumke, G. Burns, F. H. Dill, and G. Lasher, "Stimulated emission of radiation from GaAs p-n junctions," *Applied Physics Letters*, vol. 1, no. 3, p. 62, 1962.
[94] M. Abe, "Silica-based waveguide devices for photonic networks," *Journal of the Ceramic Society of Japan*, vol. 116, no. 1358, pp. 1063–1070, 2008.
[95] E. Wooten, K. Kissa, A. Yi-Yan, E. Murphy, D. Lafaw, P. Hallemeier, D. Maack, D. Attanasio, D. Fritz, G. McBrien, and D. Bossi, "A review of lithium niobate modulators for fiber-optic communications systems," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 6, pp. 69–82, Jan. 2000.
[96] R. Rashed, "A pioneer in anaclastics: Ibn Sahl on burning mirrors and lenses," *Isis*, vol. 81, no. 3, pp. 464–491, 1990.
[97] E. Hecht, *Optics*. Addison Wesley, 1998.
[98] A. Yariv, *Optical Electronics in Modern Communications*. Oxford University Press, fifth ed., 1997.
[99] C. Henry and B. Verbeek, "Solution of the scalar wave equation for arbitrarily shaped dielectric waveguides by two-dimensional Fourier analysis," *Journal of Lightwave Technology*, vol. 7, no. 2, pp. 308–313, 1989.

[100] S. Hewlett and C. Kurtsiefer, "Fourier decomposition method applied to mapped infinite domains: scalar analysis of dielectric waveguides down to modal cutoff," *Journal of Lightwave Technology*, vol. 13, pp. 375–383, Mar. 1995.

[101] A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. D. Joannopoulos, and S. G. Johnson, "Meep: A flexible free-software package for electromagnetic simulations by the FDTD method," *Computer Physics Communications*, vol. 181, pp. 687–702, Mar. 2010.

[102] S. G. Johnson and J. D. Joannopoulos, "Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis," *Optics Express*, vol. 8, p. 173, Jan. 2001.

[103] W. van Etten and J. van der Plaats, *Fundamentals of optical fiber communications*. Prentice Hall, 1991.

[104] F. Adikan, *Direct UV-Written Waveguide Devices*. PhD thesis, University of Southampton, 2007.

[105] S. Nolte, M. Will, J. Burghoff, and a. Tuennermann, "Femtosecond waveguide writing: a new avenue to three-dimensional integrated optics," *Applied Physics A: Materials Science & Processing*, vol. 77, pp. 109–111, June 2003.

[106] M. Kohnen, P. G. Petrov, R. A. Nyman, and E. A. Hinds, "Minimally destructive detection of magnetically trapped atoms using frequency-synthesized light," *New Journal of Physics*, vol. 13, p. 085006, Aug. 2011.

[107] P. Tandon and H. Boek, "Experimental and theoretical studies of flame hydrolysis deposition process for making glasses for optical planar devices," *Journal of Non-Crystalline Solids*, vol. 317, pp. 275–289, Mar. 2003.

[108] M. Kristensen, "Bragg Gratings White Paper." http://www.ibsenphotonics.com/phasemasks/tutorials/bragggratingswhitepaper, 2005.

[109] M. Svalgaard, C. Poulsen, A. Bjarklev, and O. Poulsen, "Direct UV writing of buried singlemode channel waveguides in Ge-doped silica films," *Electronics Letters*, vol. 30, no. 17, pp. 1401–1403, 1994.

[110] G. Emmerson, S. Watts, C. Gawith, V. Albanis, M. Ibsen, R. Williams, and P. G. R. Smith, "Fabrication of directly UV-written channel waveguides with simultaneously defined integral Bragg gratings," *Electronics Letters*, vol. 38, no. 24, p. 1531, 2002.

[111] D. Kundys, J. C. Gates, S. Dasgupta, and P. G. R. Smith, "Use of cross-couplers to decrease size of UV written photonic circuits," *Photonics Technology Letters, IEEE*, vol. 21, no. 13, pp. 947–949, 2009.

[112] F. Adikan, C. Gawith, P. G. R. Smith, I. J. G. Sparrow, G. Emmerson, C. Riziotis, and H. Ahmad, "Design and demonstration of direct UV-written small angle X couplers in silica-on-silicon for broadband operation.," *Applied optics*, vol. 45, pp. 6113–8, Aug. 2006.

[113] I. J. G. Sparrow, G. Emmerson, C. Gawith, and P. G. R. Smith, "Planar waveguide hygrometer and state sensor demonstrating supercooled water recognition," *Sensors and Actuators B: Chemical*, vol. 107, pp. 856–860, June 2005.

[114] H. Ebendorff-Heidepriem, "Laser writing of waveguides in photosensitive glasses," *Optical Materials*, vol. 25, pp. 109–115, Mar. 2004.

[115] H. L. Rogers, S. Ambran, C. Holmes, P. G. R. Smith, and J. C. Gates, "In situ loss measurement of direct UV-written waveguides using integrated Bragg gratings.," *Optics letters*, vol. 35, pp. 2849–51, Sept. 2010.

[116] G. M. Hale and M. R. Querry, "Optical Constants of Water in the 200-nm to 200-microm Wavelength Region.," *Applied optics*, vol. 12, pp. 555–63, Mar. 1973.

[117] K. M. Davis, K. Miura, N. Sugimoto, and K. Hirao, "Writing waveguides in glass with a femtosecond laser.," *Optics letters*, vol. 21, pp. 1729–31, Dec. 1996.

[118] C. Dongre, R. Dekker, H. J. W. M. Hoekstra, and D. Nolli, "Characterization of Femtosecond Laser Written Waveguides for Integrated Biochemical Sensing," in *Proc. IEEE, LEOS Benelux 2007*, (Brussels), pp. 1–4, 2007.

[119] M. Ams, G. D. Marshall, P. Dekker, J. Piper, and M. J. Withford, "Ultrafast laser written

active devices," *Laser & Photonics Review*, vol. 3, pp. 535–544, Nov. 2009.
[120] K. C. Vishnubhatla, S. V. Rao, R. S. S. Kumar, R. Osellame, S. N. B. Bhaktha, S. Turrell, A. Chiappini, A. Chiasera, M. Ferrari, M. Mattarelli, M. Montagna, R. Ramponi, G. C. Righini, and D. N. Rao, "Femtosecond laser direct writing of gratings and waveguides in high quantum efficiency erbium-doped Baccarat glass," *Journal of Physics D: Applied Physics*, vol. 42, p. 205106, Oct. 2009.
[121] J. Thomas, M. Heinrich, P. Zeil, V. Hilbert, K. Rademaker, R. Riedel, S. Ringleb, C. Dubs, J.-P. Ruske, S. Nolte, and A. Tünnermann, "Laser direct writing: Enabling monolithic and hybrid integrated solutions on the lithium niobate platform," *Physica Status Solidi (a)*, vol. 208, pp. 276–283, Feb. 2011.
[122] L. Sansoni, F. Sciarrino, G. Vallone, P. Mataloni, A. Crespi, R. Ramponi, and R. Osellame, "Polarization Entangled State Measurement on a Chip," *Physical Review Letters*, vol. 105, pp. 1–4, Nov. 2010.
[123] A. Snyder, "Coupled-mode theory for optical fibers," *JOSA*, vol. 62, no. 11, pp. 1267–1277, 1972.
[124] S. Fleming and H. An, "Poled glasses and poled fibre devices," *Journal of the Ceramic Society of Japan*, vol. 116, no. 1358, pp. 1007–1023, 2008.
[125] J. H. Wray and J. T. Neu, "Refractive Index of Several Glasses as a Function of Wavelength and Temperature," *Journal of the Optical Society of America*, vol. 59, p. 774, June 1969.
[126] A. Politi, M. J. Cryan, J. G. Rarity, S. Yu, and J. L. O'Brien, "Silica-on-silicon waveguide quantum circuits.," *Science*, vol. 320, pp. 646–9, May 2008.
[127] J. Zhou, T. R. Ohno, and C. a. Wolden, "High-temperature stability of nichrome in reactive environments," *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, vol. 21, no. 3, p. 756, 2003.
[128] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 1999.
[129] M. Trupke, *Microcavities for atom chips*. PhD thesis, Imperial College London, 2007.
[130] H. Benisty, R. Stanley, and M. Mayer, "Method of source terms for dipole emission modification in modes of arbitrary planar structures," *Journal of the Optical Society of America A*, vol. 15, no. 5, p. 1192, 1998.
[131] R. Ramponi, "Femtosecond laser micromachining for optofluidic devices. Imperial College Seminar, 2010," tech. rep., 2010.
[132] C. Technologies, "Silica Waveguide Fabrication Capability Fact Sheet," tech. rep., 2011. http://www.ciphotonics.com/download/factsheet/fs-silica-waveguide-fabrication.pdf.
[133] A. Wilson, H. Takahashi, A. Riley-Watson, F. Orucevic, P. Blythe, A. Mortensen, D. R. Crick, N. Seymour-Smith, E. Brama, M. Keller, and W. Lange, "Fiber-coupled single ion as an efficient quantum light source," *arXiv*, Jan. 2011, 1101.5877.
[134] M. Steiner, H. M. Meyer, C. Deutsch, J. Reichel, and M. Köhl, "Single Ion Coupled to an Optical Fiber Cavity," *Physical Review Letters*, vol. 110, p. 043003, Jan. 2013, 1211.0050.
[135] A. Muller, E. B. Flagg, M. Metcalfe, J. Lawall, and G. S. Solomon, "Coupling an epitaxial quantum dot to a fiber-based external-mirror microcavity," *Applied Physics Letters*, vol. 95, no. 17, p. 173101, 2009.
[136] V. Kaajakari, *Practical MEMS*. Small Gear Publishing, 2009.
[137] N. Miladinovic, F. Hasan, N. Chisholm, I. Linnington, E. A. Hinds, and D. H. O'Dell, "Adiabatic transfer of light in a double cavity and the optical Landau-Zener problem," *Physical Review A*, vol. 84, p. 043822, Oct. 2011.
[138] R. J. Lang and A. Yariv, "Coupling coefficients for coupled-cavity lasers," *Quantum Electronics, IEEE Journal of*, vol. 23, no. 3, pp. 287–289, 1987.
[139] J. Poon, P. Chak, J. Choi, and A. Yariv, "Slowing light with Fabry – Perot resonator arrays," *JOSA B*, vol. 24, no. 11, 2007.
[140] S. Mookherjea and A. Yariv, "Coupled resonator optical waveguides," *Selected Topics in*

Quantum Electronics, IEEE Journal of, vol. 8, no. 3, pp. 448–456, 2002.

[141] G. Lepert, M. Trupke, E. A. Hinds, H. L. Rogers, J. C. Gates, and P. G. R. Smith, "Demonstration of UV-written waveguides, Bragg gratings and cavities at 780 nm, and an original experimental measurement of group delay," *Optics Express*, vol. 19, p. 24933, Nov. 2011.

[142] T. Erdogan, "Fiber grating spectra," *Lightwave Technology, Journal of*, vol. 15, no. 8, pp. 1277–1294, 1997.

[143] J. F. Bauters, M. Heck, and D. John, "Ultra-Low-Loss (<0.1 dB/m) Planar Silica Waveguide Technology," in *ECOC 2011*, 2011.

[144] F. Pampaloni and J. Enderlein, "Gaussian, Hermite-Gaussian, and Laguerre-Gaussian beams: A primer," *arXiv*, Oct. 2004, 0410021.

[145] H. Kogelnik and T. Li, "Laser Beams and Resonators," *Applied Optics*, vol. 5, p. 1550, Oct. 1966.

[146] D. Marcuse, "Loss analysis of single-mode fibre splices," *Bell Syst. Tech J.*, vol. 56, no. 5, p. 703, 1977.

[147] J. J. Sakurai, *Modern Quantum Mechanics*. Addison-Wesley, revised ed., 1994.