Quantum Memories
and
Monolithic Resonators

TWO APPROACHES TO TUNABLE OPTICAL ELEMENTS
FOR QUANTUM INFORMATION PROCESSING

Geoff Campbell

A thesis presented for the degree of
Doctor of Philosophy of
The Australian National University

Centre for Quantum Computation & Communication Technology
Department of Quantum Science

2 March 2015

Declaration

This thesis is an account of research undertaken between July 2010 and March 2015 at The Department of Quantum Science, Research School of Physics, Faculty of Science, The Australian National University, Canberra, Australia.

Except where acknowledged in the customary manner, the material presented in this thesis is, to the best of my knowledge, original and has not been submitted in whole or part for a degree in any university.

Geoff Campbell
2 March 2015

Acknowledgements

The work presented in this thesis wouldn't have been possible without the contributions from a long list of people. The work builds off the efforts of an amazing group of individuals and was overseen by excellent supervisors. I have been supported academically and personally by many people over the last four and a half years, all of whom deserve more thanks than I can give them in a few short paragraphs.

The research efforts were directed by my supervisors Ping Koy Lam, Ben Buchler and Boris Hage. They have all inspired me with their enthusiasm and deep knowledge of the field. Ping Koy and Ben have worked hard to make the research team a creative environment for thinking about our work in new ways and were always open to letting me try out new ideas. They have always placed the highest importance on the best interests of their students and always had an open door to discuss any difficulties that arose during my studies. I also never ceased to be amazed at Boris's capacity to build anything that was needed in the lab or just for fun. I continue to take inspiration from his ability to solve most problems before I could figure out how to even approach them.

The research that I am reporting here is also inextricably linked to the work done by my colleagues. Everything presented in this thesis is the product of a team effort by many students and the supervisors working together to explore interesting physics. The work on quantum memories was done in close collaboration with Mahdi Hosseini, Ben Sparkes, Jiao Geng, Julien Bernu, Daniel Higginbottom, Jesse Everett, Olivier Pinel, Young-Wook Cho, Quentin Glorieux and Mingtao Cao. The experiments were a result of all of these people working together and would not be possible without each of them. Mahdi Hosseini, in particular, is responsible for making Gradient Echo Memory the robust platform that it is toady and the cold-atom experiment is the product of super-human efforts by Ben Sparkes, Julien Bernu and Nick Robins. They have built an experiment that promises to yield exciting new physics for years to come. A special thanks goes to Nick Robins, Kyle Hardman and the rest of their team who provided us with access to their extensive knowledge of magneto-optical trapping and vacuum systems.

The work on monolithic resonators was a team effort with Shen Yong, Alexandre Briussel, Boris Hage, and Jiri Janousek. Boris and Jiri provided invaluable expertise to the project while Shen Yong and Alexandre became experts at the subtle arts of aligning coupling prisms and polishing the resonators. The results that we obtained with monolithic resonators are due to the patience and diligence of Shen Yong and Alex.

I would also like to thank all of my colleagues that I did not work with directly in the lab but who made it fun to be at work everyday and whose insights I came to rely on. I had many fascinating discussions with Giovanni Guccione, Seiji Armstrong, Helen Chrzanowski, Oliver Thearle, Jing-Yan Haw, Sara Hosseini, Zhao Jie, Harry Slatyer, Jeremy Gallant and Syed M Assad. My studies would have been far less fun and interesting without them. Amanda White has also been invaluable in helping navigate the bureaucracy of the university and was always ready to my mind off physics with entertaining stories when the experiments weren't working.

The biggest thanks goes to my amazing partner, Ellen. We've had a great time moving to Australia and making wonderful friendships. She has also been impossibly supportive through the stressful times of completing a PhD and I couldn't have done it without her. In addition to intangible support, Ellen has also made a substantial tangible contribution to this thesis in form of direct grammatical interventions throughout the writing process.

Finally, I would like thank my parents, who got my interest in science started when I was young, all of my friends who have made the last four and half years awesome, and all of the interesting people that I have had the opportunity to meet through my work.

Abstract

This thesis presents progress towards two complementary goals. The first part discusses our aim of building high-efficiency quantum memories for quantum optical states and exploring the potential of the memory for performing operations as well as for passive storage. The second part presents our work towards developing a compact source of quantum light by fabricating optical resonators from non-linear optical crystals.

Our high-efficiency memories operate by storing the information carried by a weak laser field in the coherence between the hyperfine-split electronic energy levels in a vapour of atomic rubidium. We analyse the performance of a three-level gradient echo memory protocol that is used to store and retrieve optical information from the atomic coherence using an off-resonant Raman transition. Using a warm rubidium vapour, we find that the memory preserves the quantum state of the input light better than any classical means of storage could achieve. We also present results showing that trapped, cold atomic vapours provide an increased storage lifetime and analyse the contribution that four-wave-mixing makes to the retrieved light. Both the warm- and cold-atom implementations of the memory are shown to store and recall light with over 80% efficiency.

We also present a number of extensions to the basic memory protocol. We demonstrate that the light-atom coupling can be controlled by the bright coupling field that is used to complete the Raman transition. This control is used to perform a beam-splitting operations between two optical pulses that are separated in time. Using frequency-multiplexed states, we show that beam-splitting operations are also possible between optical modes that are separated in frequency. This operation can be extended to perform arbitrary operations by using a series of memories. Finally, we present a high-fidelity memory for dual-rail q-bits encoded across two frequency modes.

The resonators that were fabricated to explore compact sources of quantum light demonstrated many of the properties that are required to obtain the stable operation of a quantum source. We demonstrated that non-linear optical crystals can be machined into monolithic resonators that use total internal reflections to confine light. Two geometries were tested, whispering gallery resonators and square travelling wave cavities. Because both types of resonators use total internal reflection instead of dielectric coatings, they feature very low loss rates. Lithium niobate resonators of both types exhibited cavity decay rates that were dominated by bulk absorption in the crystal. The resonators were shown to be frequency tunable using an applied voltage and could be coupled to free-space modes with a tunable coupling rate. We present preliminary results of non-linear frequency conversion in the resonators.

We also investigated calcium fluoride and magnesium fluoride resonators for use as high-finesse reference cavities. We report a finesse of ≈ 5000 in a calcium fluoride whispering gallery resonator. A square calcium fluoride monolithic resonator was also produce to have a low insertion loss when coupled to a Gaussian free-space mode. We demonstrate a power coupling efficiency of 94% to that resonator. We also present a proposal to use a square lithium niobate resonator as a tunable coupler for a high-finesse calcium fluoride whispering gallery resonator. The coupled resonator system could be used as a compact high-efficiency optical quantum memory.

Contents

List of Figures

List of Tables

Introduction

CHAPTER 1

Background and motivation

Within efforts to engineer quantum computers, or at least the basic building blocks of such devices, optics occupies an interesting space. Single photons are relatively easy to produce, survive interactions with their surroundings and travel quickly. The ease with which single photons and entangled photons can be produced has resulted in optics being an excellent platform to test fundamental properties of quantum mechanics. Additionally, the fact that photons can be transmitted over long distances without being lost or altered secures optics as being the only viable means to transmit quantum information between distant locations.

The work in this thesis is focused on the long-term goal of developing technologies and techniques to further quantum optics as a platform for manipulating and communicating quantum information. Beyond that broad outline, the work is divided into two main parts. The first part covers a collection of experiments and ideas about memories that are capable of storing and retrieving single photons or other optical states without disturbing their quantum state. These memories rely on large numbers of atoms to absorb and re-emit light in a controllable manner. The second part of the thesis covers the development of compact optical resonators with the goal of using them to generate quantum states of light, namely squeezed light or entangled pairs of photons. These parts reflect two important aspects of developing a network for communicating quantum information, the ability to create quantum light and the ability to manipulate it as it travels through the network.

The capacity to store and retrieve quantum information that is encoded onto light is a necessity for long-distance quantum communication. Quantum states are extremely fragile to any disturbance, such as loss in a fibre, and simply cannot be transmitted over long distances through optical fibre networks. The standard technique of amplifying signals to counteract loss will introduce unavoidable noise to a quantum signal and provide no advantage. The solution is to teleport information directly from the transmitter to the receiver and thereby bypass loss along the way. This however, requires distributing the two parts of an entangled state to the sender and receiver through the same lossy channels that will the destroy the signal. To avoid destroying the entangled resource as well, the long link between the sender and receiver can be broken into shorter segments. If entangled photons can survive transmission through each short segment and be stored at either end in memories, the entanglement can be extended across all segments by measurements of the stored photons at each node. In this manner, the endpoints of the network become entangled and a signal can be teleported from end to end with no loss.

Of course, storing and manipulating optical quantum information is only part of the picture. The quantum states must be created before they can be manipulated and this is the focus of the second part of this thesis. Optical quantum state engineering is more successful than efforts to engineer memories and repeater networks. Some of the earliest demonstrations of the nature of quantum states were made through by the preparation of entangled pairs of photons. The precess is straightforward, it is possible to double the frequency of light by using non-linear optical elements and the process is reversible. The reverse process results in high-energy photons splitting into two photons of half the energy. If these lower energy photons cannot be distinguished, they are entangled. This process is commonly used in quantum optics, to the point that it can be done routinely in many labs. The project to develop compact resonators aims to miniaturise and simplify the optical systems needed to accomplish this such that quantum light is easily available for more complex protocols.

While the two parts of this thesis discuss very different optical systems with different aims, they have the same motivation: to explore the interesting physics behind systems that have potential applications in optical quantum computation and communication.

CHAPTER 2

Publications

Publications resulting from work contained in this thesis

1. Hosseini, M., Sparkes, B. M., Campbell, G., Lam, P. K. & Buchler, B. C. High efficiency coherent optical memory with warm rubidium vapour. *Nat. Commun.* **2** (Feb. 2011).
2. Hosseini, M., Campbell, G., Sparkes, B. M., Lam, P. K. & Buchler, B. C. Unconditional room temperature quantum memory. *Nat. Phys., DOI: 10.1038/NPHYS2021.* doi:`DOI:10.1038/NPHYS2021` (2011).
3. Campbell, G., Hosseini, M., Sparkes, B. M., Lam, P. K. & Buchler, B. C. Time- and frequency-domain polariton interference. *New J. of Phys.* **14,** 033022 (2012).
4. Hosseini, M., Sparkes, B. M., Campbell, G. T., Lam, P. K. & Buchler, B. C. Storage and manipulation of light using a Raman gradient-echo process. *Journal of Physics B: Atomic, Molecular and Optical Physics* **45,** 124004 (2012).
5. Campbell, G. T. *et al.* Configurable Unitary Transformations and Linear Logic Gates Using Quantum Memories. *Phys. Rev. Lett.* **113,** 063601 (6 Aug. 2014).
6. Geng, J. *et al.* Electromagnetically induced transparency and four-wave mixing in a cold atomic ensemble with large optical depth. *New Journal of Physics* **16,** 113053 (2014).

Publications authored during my PhD not covered in this thesis

1. Sparkes, B. M. *et al.* Precision Spectral Manipulation: A Demonstration Using a Coherent Optical Memory. *Phys. Rev. X* **2,** 021011 (2 June 2012).
2. Weng, W. *et al.* Nano-Kelvin Thermometry and Temperature Control: Beyond the Thermal Noise Limit. *Phys. Rev. Lett.* **112,** 160801 (16 Apr. 2014).
3. Campbell, G., Hage, B., Buchler, B. & Lam, P. K. Generation of high-order optical vortices using directly machined spiral phase mirrors. *Appl. Opt.* **51,** 873–876 (Mar. 2012).

4. Shen, Y. *et al.* Generation and interferometric analysis of high charge optical vortices. *Journal of Optics* **15,** 044005 (2013).
5. Genevaux, P. *et al.* *3 Modes transmission using hybrid separation with high mode selectivity and low losses spatial mode multiplexer* in (IEEE, 2014), 1–3. doi:`10.1109/ECOC.2014.6964119`.

CHAPTER 3

Quantum light

The goal of the experiments discussed in this thesis is to prepare and manipulate the quantum state of light. The first task must therefore be to describe what exactly is meant by a quantum state of light. The definition of what distinguishes a quantum state from a classical one is complex and occasionally contentious but could generally be described as a state that cannot be described without discussing the discrete nature of light, for example, a single photon.

For the purposes of this thesis however, I use the term quantum state of light loosely. The experiments that were performed do not involve light that satisfies the criterion of being non-classical, however, the notion that light is quantum is still central to the motivation for the work. With that mind, it is worthwhile providing a brief description of light as a quantum phenomenon and to discuss how the states used in this work relate to quantum states. When discussing the quantum state of light, the concept of a mode is often introduced. A mode of light is a localised region in space and time that is a valid solution of the Maxwell equations. The mode is described by a weight function that corresponds to the classical field amplitudes for the particular solution. The quantum state of the mode can be described and, if the mode can be shone onto a detector, the number of photons that it contains can be measured.

This idea of a mode is extremely important in quantum optics and a good starting point for formulating how a mode is described is to discuss how light propagates, the equations that govern its dynamics and how it behaves as a quantum field. This is done in numerous textbooks [b1–b4], nonetheless I will repeat it here because it is necessary to adequately explain what is meant by an optical mode, a concept that is pervasive in the rest of the thesis.

3.1 The Maxwell Equations

The propagation of electromagnetic fields is described by the Maxwell equations [b5]:

$$\nabla \cdot \mathbf{D} = \rho, \tag{3.1a}$$

$$\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t}, \tag{3.1b}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{3.1c}$$

$$\nabla \times \mathbf{E} + \frac{\partial \mathbf{B}}{\partial t} = 0. \tag{3.1d}$$

These coupled equations, with source terms for the charge density, ρ, and current density, $\mathbf{J}$, describe the distribution and propagation of the electric, $\mathbf{E}$, and magnetic, $\mathbf{H}$, fields and the electric displacement, $\mathbf{D} = \epsilon_0 \epsilon \mathbf{E}$, and magnetic induction, $\mathbf{B} = \mu_0 \mu \mathbf{H}$. For light, however, it is more practical to reduce the set of equations to two second-order differential equations. This can be done with the introduction of a scalar potential, Φ, and vector potential, $\mathbf{A}$. In terms of these potentials, the Maxwell equations can be written:

If no source terms are present, it is convenient to use the Coulomb gauge, $\nabla \cdot \mathbf{A} = 0$, so that $\Phi = 0$. In this case, the only remaining equation is

$$\nabla^2 \mathbf{A} - \frac{1}{c^2} \frac{\partial^2 \mathbf{A}}{\partial t^2} = 0, \tag{3.2}$$

and the electric and magnetic fields are given by

$$\mathbf{E} = -\frac{\partial \mathbf{A}}{\partial t}, \tag{3.3a}$$

$$\mathbf{B} = \nabla \times \mathbf{A}. \tag{3.3b}$$

We now have a simple description of the propagation of light in free space given in terms of the vector potential eq. (3.3a). The propagation of light can be described by finding solutions for eq. (3.2) and, through equations eq. (3.3a) and eq. (3.3b), the real fields $\mathbf{E}$ and $\mathbf{B}$ can be calculated.

If we assume a solution to 3.2 that has time-dependence of the form $e^{-i\omega t}$, we can write it as

$$(\nabla^2 + k^2)\mathbf{A} = 0, \tag{3.4}$$

where $k = \frac{\omega}{c}$. This is the Helmholtz wave equation, which has a solution of the form,

$$\mathbf{u_k}(\mathbf{r}, t) = a\mathbf{e}^{(\lambda)} e^{i\mathbf{k}\cdot\mathbf{r} - i\omega_k t}, \tag{3.5}$$

where $\mathbf{e}^{(\lambda)}$ is the polarisation vector with index $\lambda = 1, 2$ for two transverse, orthogonal polarisations.

Equation (3.5) describes a plane wave of infinite spatial extent propagating in the direction of $\mathbf{k}$; not a solution that is physical, but one that is quite useful due to a consequence of the linearity of the Maxwell equations eqs. (3.1). The superposition principle applies to the electromagnetic fields, therefore a linear superposition of solutions is itself a solution. From this linearity, we can construct an arbitrary solution for the vector potential from the fact that with a given volume V the solutions $\mathbf{u_k}$ form an orthonormal basis,

$$\int_V \mathbf{u}^*_{\mathbf{k'}}(\mathbf{r}) \mathbf{u_k}(\mathbf{r}) d\mathbf{r} = \delta_{k',k} \tag{3.6}$$

if we adopt the normalisation

$$\mathbf{u}_{\mathbf{k}}(\mathbf{r}, t) = V^{-1/2}\mathbf{e}^{(\lambda)}e^{i\mathbf{k}\cdot\mathbf{r}-i\omega_k t}. \tag{3.7}$$

This applies to fields that vanish at the boundaries, or, equivalently, for volumes enclosed within reflecting boundary conditions. The fact that a general solution to eq. (3.2) in free-space can be decomposed into plane-waves has useful implications for analysing propagating optical fields. Any results that apply to a plane wave, a one-dimensional problem, also apply to a general three-dimensional field. Furthermore, under the paraxial approximation, which will be discussed in section 3.5, it allows a simple description for optical fields which propagate as narrow beams, such as the optical fields generated by lasers.

3.2 Quantum states

In quantum mechanics, the state of a physical system is represented by a vector. This vector exists in a vector space that includes every possible physical configuration of the system. The evolution of physical systems that are governed by quantum mechanics is therefore a problem that is described in the language of linear algebra. The representation used to describe the state vectors is usually Dirac notation. In this formalism, a vector $\vec{v}$ is written as $|v\rangle$. A state vector $|\psi\rangle$ can always be decomposed into basis states

$$|\psi\rangle = \sum_i \psi_i |u_i\rangle \tag{3.8}$$

where the vectors $\{|u_i\rangle\}$ form an orthonormal basis. The conjugate transpose of the vector is denoted $|\psi\rangle^\dagger \equiv \langle\psi|$ and the inner product between two vectors $|\psi\rangle$ and $|\phi\rangle$ is

$$\langle\psi|\phi\rangle \equiv \sum_i \psi_i^* \phi_i. \tag{3.9}$$

The state vector of a physical system is normalised such the $\langle\psi|\psi\rangle \equiv \sum_i |\psi_i|^2 = 1$.

Any measurement that can be made on a quantum system, for example, the position, mass or charge of a particle, is associated with a Hermitian operator acting in the same vector space of the system. The notation used for operators is to place a hat on them, for example an operator acting in a space of dimension N is

$$\hat{A} = \sum_{i,j}^{N} a_{i,j} |a_i\rangle\langle a_j| \tag{3.10}$$

and common shorthand is

$$\hat{A}_{i,j} \equiv a_{i,j} |a_i\rangle\langle a_j|. \tag{3.11}$$

Using spectral decomposition, any observable can be diagonalised and written as $\hat{O} = \sum_i o_i |o_i\rangle\langle o_i|$ where $\{|o_i\rangle\}$ are the eigenvectors of the operator. The observable acting on a state will give an expectation value

$$\langle\hat{O}\rangle_\psi \equiv \langle\psi|\hat{O}|\psi\rangle = \sum_i |\langle o_i|\psi\rangle|^2 \tag{3.12}$$

The value $|\langle o_i|\psi\rangle|^2$ is interpreted as the probability that a measurement of the system $|\psi\rangle$ will find that it is in the state $|o_i\rangle$.

3.2.1 The Schrödinger equation

One observable that is applicable to all quantum systems is that associated with the energy of the system. This is known as the Hamiltonian and determines the dynamics of a quantum system. The Hamiltonian, written in its eigenbasis $\{|E_i\rangle\}$, is

$$\hat{H} = \sum_i^N E_i |E_i\rangle\langle E_i| \tag{3.13}$$

where the eigenvalues E_i have units of energy. The fundamental postulate of quantum mechanics, the Schrödinger equation, is written in terms of the Hamiltonian:

$$i\hbar\partial_t|\psi(t)\rangle = \hat{H}(t)|\psi(t)\rangle. \tag{3.14}$$

For a constant Hamiltonian, eq. (3.14) has the solution

$$|\psi(t)\rangle = e^{-i\hat{H}t/\hbar}\psi(0). \tag{3.15}$$

3.3 Quantum fields

Study of the classical fields that are governed by the Maxwell equations necessarily involves measurement of those fields. If we now suppose that the quantities which we measure are, in fact, quantum phenomena, measurements of those fields must relate to the expectation values of the corresponding field operators. These field operators are an expansion of the field over a complete set of orthonormal functions. The solutions, eq. (3.7), to eq. (3.4) are just such an expansion and we can define operators for the vector potential:

$$\hat{\mathbf{A}}(\mathbf{r},t) = \sqrt{\frac{\hbar}{2\omega V\epsilon_0}}\sum_k (\mathbf{A}_k(\mathbf{r},t)\hat{a}_k + \mathbf{A}_k^*(\mathbf{r},t)\hat{a}_k^\dagger), \tag{3.16}$$

where the coefficients, $\mathbf{A}_k(\mathbf{r},t)$, take the form of eq. (3.5). The second term is the complex conjugate of the first and is present because $\mathbf{A}$ is a real quantity. The operators are subject to the bosonic commutation relations

$$[\hat{a}_k, \hat{a}_{k'}] = 0, \tag{3.17a}$$

$$[\hat{a}_k^\dagger, \hat{a}_{k'}^\dagger] = 0, \tag{3.17b}$$

$$[\hat{a}_k, \hat{a}_{k'}^\dagger] = \delta_{k,k'}. \tag{3.17c}$$

From the linearity of the Maxwell equations we know that, in the Heisenberg picture the field operators themselves must obey the same dynamics that govern the classical fields [b2]. We therefore know how $\hat{\mathbf{A}}$ propagates and from it we can find $\hat{\mathbf{E}}$, $\hat{\mathbf{H}}$, $\hat{\mathbf{D}}$ and $\hat{\mathbf{B}}$. The operator for the electric field, for example, is found from eqs. (3.3a) and (3.16):

$$\hat{\mathbf{E}}(\mathbf{r},t) = i\sqrt{\frac{\hbar}{2V\epsilon_0}}\sum_k \sqrt{\omega_k}(\mathbf{A}_k(\mathbf{r},t)\hat{a}_k + \mathbf{A}_k^*(\mathbf{r},t)\hat{a}_k^\dagger). \tag{3.18}$$

The operators apply for a quantisation volume V that is enclosed by reflecting boundary conditions; however, the procedure is valid even as the volume becomes arbitrarily large and the the mode density arbitrarily high. Any physical vector potential can be represented as a superposition of operators and the quantisation remains valid. We can take the electric field operator out of its box by considering eq. (3.18) in the limit of an arbitrarily long box. Replacing the sum over mode numbers with an integral yields

$$\hat{\mathbf{E}}(\mathbf{r},t) = i\sqrt{\frac{\hbar\omega_0}{2V\epsilon_0}}\int_{\omega}(\mathbf{A}_\omega(\mathbf{r},t)\hat{a}_\omega + \mathbf{A}^*_\omega(\mathbf{r},t)\hat{a}^\dagger_\omega). \tag{3.19}$$

where ω_0 is a carrier frequency and the integral is performed over a bandwidth that contains the spectrum of interest.

3.3.1 Field energy and the number operator

The Hamiltonian for the quantum optical field is the operator corresponding to the energy contained in the field. Substituting the field operators, $\hat{\mathbf{E}}$ and $\hat{\mathbf{B}}$, into the classical expression for the energy contained in an electromagnetic field provides the Hamiltonian operator. The normalization of the vector potential operator, eq. (3.16), ensures the convenient form,

$$\hat{H} = \sum_k \hbar\omega\left(\hat{a}^\dagger_k\hat{a}_k + \frac{1}{2}\right). \tag{3.20}$$

This can be identified as the Hamiltonian for a quantum harmonic oscillator and he operator

$$\hat{n} = \hat{a}^\dagger_k\hat{a}_k \tag{3.21}$$

is the number operator for the number of photons in the mode $\hat{a}_k$, each with an energy of $\hbar\omega$. The operators $\hat{a}^\dagger$ and $\hat{a}$ can also be identified as the raising and lowering operators respectively. The application of these to a quantum mode adds or removes one photon.

3.3.2 Optical phase-space

The electric field is a real quantity and we can express the electric field operator in terms of real quantities as

$$\hat{E} \propto \hat{q}\cos(\omega t - \mathbf{k}\cdot\mathbf{r}) + \hat{p}i\sin(\omega t - \mathbf{k}\cdot\mathbf{r}). \tag{3.22}$$

From eq. (3.19), the definitions of $\hat{q}$ and $\hat{p}$ are

$$\hat{q} = \frac{1}{\sqrt{2}}\left(\hat{a}^\dagger + \hat{a}\right) \tag{3.23a}$$

$$\hat{p} = -\frac{i}{\sqrt{2}}\left(\hat{a}^\dagger - \hat{a}\right). \tag{3.23b}$$

These are referred to as the field quadratures, in analogy to the in-phase and out-of-phase quadratures in signal processing.

3.3.3 Optical states

We can define states that a given optical mode can occupy. The most straight forward is the Fock state, the eigenstate of the number operator. These states correspond to a mode that contains a well defined number of photons. We use the notation $\{|0\rangle, |1\rangle, |2\rangle ... |n\rangle\}$

to denote Fock states that contain $\{0, 1, 2, ...n\}$ photons. The action of the raising and lowering operators on Fock states gives

$$\hat{a}|n\rangle = \sqrt{n}|n-1\rangle \qquad \hat{a}^\dagger|n\rangle = \sqrt{n+1}|n+1\rangle. \tag{3.24}$$

The only Fock state that can be produced trivially in the lab is the vacuum state $|0\rangle$ which contains no photons.

The most common pure state, other than $|0\rangle$, that is encountered experimentally is the coherent state, denoted $|\alpha\rangle$ where α is the amplitude of the state. It resembles a classical oscillator, with the amplitude of the state evolving in phase space at the optical frequency. Coherent states consist of phase-coherent photons with uncorrelated arrival times and are produced by ideal lasers (those that produce the lowest noise permitted by the uncertainty principle). The coherent state is the eigenstate of the lowering operator

$$\hat{a}|\alpha\rangle = \alpha|\alpha\rangle \tag{3.25}$$

and can be decomposed into Fock states according to

$$|\alpha\rangle = e^{-\frac{|\alpha|^2}{2}} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}}|n\rangle. \tag{3.26}$$

It is also a minimum uncertainty state with equal noise when projected onto any quadrature.

If the noise can be reduced in one quadrature, we have a squeezed coherent state. The squeezed vacuum state is a special case of the squeezed coherent state that has no mean amplitude. In a squeezed vacuum state, the photon arrival times are correlated, specifically, they must always arrive in pairs. The squeezed vacuum state can therefore be created by down-conversion, the process of converting one photon into two photons of lower energy. A much more detailed description of quantum states can be found in [b2]

3.4 The propagation of quantum light

The propagation of the quantum fields is governed by the Helmholtz wave equation eq. (3.4). Directly solving the wave equation, however, is generally intractable on the scales of an optical table or laser cavity because the modes vary on very small length- and time-scales. The usual approach is to consider how slowly-varying envelopes of the modes vary under assumptions that are usually satisfied. Often, the spatial and temporal modes of an optical pulse are considered separately. The spatial mode is derived assuming a continuous monochromatic beam while the temporal evolution is considered only for one dimension by assuming a plane wave propagating along the axis of the beam.

3.4.1 Pulse propagation

Let us first examine the propagation of a plane-wave pulse of light. The temporal evolution the optical pulse can be found from the macroscopic Maxwell equations for fields eqs. (3.1). For a field propagating in a dielectric material, the Helmholtz wave equation, eq. (3.4), becomes

$$\nabla^2\mathbf{E} - \frac{1}{c^2}\frac{\partial^2\mathbf{E}}{\partial t^2} = \mu_0\frac{\partial^2\mathbf{P}}{\partial t^2} \tag{3.27}$$

where $\mathbf{P}$ is the polarization of the dielectric. The plane wave fields can be described by

$$\mathbf{E}(\boldsymbol{r},t) = \hat{e}\mathcal{E}(z,t)e^{i(kz-\omega t)} + \hat{e}\mathcal{E}^*(z,t)e^{-i(kz-\omega t)} \tag{3.28a}$$
$$\mathbf{P}(\boldsymbol{r},t) = \hat{e}\mathcal{P}(z,t)e^{i(kz-\omega t)} + \hat{e}\mathcal{P}^*(z,t)e^{i(kz-\omega t)} \tag{3.28b}$$

where $\mathcal{E}$ and $\mathcal{P}$ are slowly varying complex envelopes and $\hat{e}$ is the polarization unit vector. Substituting eqs. (3.28) into eq. (3.27) and using the approximations

$$\left|\frac{\partial\mathcal{E}}{\partial t}\right| \ll \omega|\mathcal{E}|, \left|\frac{\partial\mathcal{E}}{\partial z}\right| \ll k|\mathcal{E}|, \tag{3.29a}$$
$$\left|\frac{\partial\mathcal{P}}{\partial t}\right| \ll \omega|\mathcal{P}|, \left|\frac{\partial\mathcal{P}}{\partial z}\right| \ll k|\mathcal{P}| \tag{3.29b}$$

we obtain

$$\left(\frac{1}{c}\partial_t + \partial_z\right)\mathcal{E} = i\frac{k}{2\epsilon_0}\mathcal{P}. \tag{3.30}$$

A common simplification of this equation is to perform a coordinate transformation to a reference frame that is moving at the speed of light [b4]

$$\xi = z, \quad \tau = t - z/c. \tag{3.31}$$

The derivatives of the transformed coordinates are

$$\frac{\partial}{\partial\xi} = \frac{\partial z}{\partial\xi}\frac{\partial}{\partial z} + \frac{\partial t}{\partial\xi}\frac{\partial}{\partial t} = \frac{\partial}{\partial z} + \frac{1}{c}\frac{\partial}{\partial t} \tag{3.32a}$$
$$\frac{\partial}{\partial\tau} = \frac{\partial z}{\partial\tau}\frac{\partial}{\partial z} + \frac{\partial t}{\partial\tau}\frac{\partial}{\partial t} = \frac{\partial}{\partial t} \tag{3.32b}$$

which yields a simple equation for the electric field envelope

$$\frac{\partial}{\partial\xi}\mathcal{E}(\xi) = i\frac{k}{2\epsilon_0}\mathcal{P}(\xi). \tag{3.33}$$

3.4.2 Beams in the paraxial approximation

With a simple equation for the temporal evolution of a plane-wave, we next move on to finding the spatial structure of an optical mode propagating in free space. We can use any set of orthonormal functions that satisfy the Maxwell equations to form a basis into which we can decompose $\hat{\mathbf{A}}$. A set of plane waves is a good choice for a box-like quantisation volume; however, a box-like region of interest is not a situation that arises frequently in the lab. It is, therefore, more helpful to decompose the modes into a set of functions that is related to the propagation of laser beams. The usual basis functions that are used are the Hermite-Gauss modes. These are a set of solutions that are shape-preserving under the action of diffraction; in other words, they are eigenmodes of free-space propagation.

We start by applying the paraxial approximation to the Helmholtz wave equation eq. (3.4). This approximation makes the assumption that the variation envelope of the beam along the direction of propagation is small on the scale of the wavelength. This is equivalent to assuming that the angle between the wave vector and the optical axis is small, hence the term paraxial approximation. Consider an optical beam with a particular polarisation propagating in z. Its amplitude is described by the scalar field $\tilde{A}$ of the form

[b6]

$$\tilde{A}(x,y,z) \equiv \tilde{u}(x,y,z)e^{-ikz}, \tag{3.34}$$

where $\tilde{u}$ is an envelope. Applying this to eq. (3.4) we get

$$\frac{\partial^2 \tilde{u}}{\partial x^2} + \frac{\partial^2 \tilde{u}}{\partial y^2} + \frac{\partial^2 \tilde{u}}{\partial z^2} - 2ik\frac{\partial \tilde{u}}{\partial z} = 0. \tag{3.35}$$

Provided that the envelope varies slowly in z,

$$|\frac{\partial^2 \tilde{u}}{\partial z^2}| \ll |2k\frac{\partial \tilde{u}}{\partial z}|, \tag{3.36}$$

this term can be dropped, and the wave equation rewritten,

$$\nabla^2 \tilde{u}(\mathbf{s},z) - 2ik\frac{\partial \tilde{u}(\mathbf{s},z)}{\partial z} = 0, \tag{3.37}$$

where $\mathbf{s} \equiv (x,y)$.

The function

$$\tilde{u}(x,y,z) = \frac{1}{\tilde{q}(z)} \times \exp\left(-ik\frac{x^2+y^2}{2\tilde{q}(z)}\right) \tag{3.38}$$

is a solution of eq. (3.37) given in terms of the complex beam parameter, $\tilde{q}(z) = z + iz_R$, where $z_R = \pi\omega_0/\lambda$ is the Rayleigh length for a beam with a wavelength λ and beam waist ω_0. This solution is more physically realisable in free space than a plane wave: it has finite spatial extent. Additionally, it has a Gaussian amplitude profile in every plane normal to the propagation direction and has phase-fronts with a spherical profile. These properties are relevant to the description of physical laser beams: mirrors with a spherical profile satisfy the boundary conditions for a beam with a finite spatial extent. Most lasers output a beam of this form because it is the lowest-order solution for an optical wave confined by spherical optics (mirrors and lenses) and this Gaussian beam is the basis for most experiments in quantum optics.

Equation (3.38) is not the only solution to the paraxial wave eq. (3.37). The solutions in each transverse direction can be found independently and their product forms a general solution. A set of solutions to the paraxial wave equation with one transverse direction are given in terms of the Hermite polynomials, H_n:

$$\begin{aligned}\tilde{u}_n(x,z) = &\left(\frac{2}{\pi}\right)^{1/4}\left(\frac{1}{2^n n!\omega_0}\right)^{1/2}\left(\frac{\tilde{q}_0}{\tilde{q}(z)}\right)^{1/2}\left(\frac{\tilde{q}_0}{\tilde{q}_0^*}\frac{\tilde{q}^*(z)}{\tilde{q}(z)}\right)^{n/2}\\ &\times H_n\left(\frac{\sqrt{2}x}{\omega(z)}\right)\exp\left(-ik\frac{x^2}{2\tilde{q}(z)}\right).\end{aligned} \tag{3.39a}$$

The set of equations given by eqs. (3.39) forms a complete set known as the Hermite-Gaussian solutions, satisfying the orthonormality condition,

$$\int_{-\infty}^{\infty} \tilde{u}_n^*(x,z)\tilde{u}_m(x,z)dx = \delta_{n,m}. \tag{3.40}$$

In every plane, an arbitrary amplitude function can be decomposed into superpositions of this basis set and the coefficients of this decomposition will remain constant under propagation due to the orthogonality of solutions. A general solution for both transverse directions is simply the product of the solutions in each independent transverse direction, thus any optical field which satisfies the paraxial approximation can be expanded in the form

$$\tilde{A}(x, y, z) = \sum_n \sum_m c_{n,m} \tilde{u}_n(x, z) \tilde{u}_m(y, z) \exp\left(-ikz\right). \tag{3.41}$$

We now have a mode expansion for the electromagnetic field and the corresponding field operators for a quantisation volume with spherical boundary conditions, such as an optical cavity. Optical fields that propagate outside of a cavity can still be described by the same mode expansion because any field can be decomposed into these basis functions.

We can rewrite eqs. (3.39) in terms of the beam waist, $\omega(z)$, the radius of curvature $R(z)$, and the Gouy phase $\psi(z)$,

$$\tilde{u}_n(x, z) = \left(\frac{2}{\pi}\right)^{1/4} \left(\frac{\exp\left[i(2n+1)\psi(z)\right]}{2^n n! \omega(z)}\right)^{1/2} \times H_n\left(\frac{\sqrt{2}x}{\omega(z)}\right) \exp\left(-ik\frac{x^2}{2R(z)} - \frac{x^2}{\omega^2(z)}\right).$$

where $\omega(z)$, $R(z)$, and $\psi(z)$ are given by the position and Rayleigh length:

$$\omega(z) = \omega_0 \sqrt{1 + \left(\frac{z}{z_r}\right)^2} = \left|\sqrt{\frac{\lambda}{\pi}\frac{1}{\mathrm{Im}[\tilde{q}^{-1}(z)]}}\right| \tag{3.43}$$

$$R(z) = z + \frac{z_R^2}{z} = \frac{1}{\mathrm{Re}[\tilde{q}^{-1}(z)]} \tag{3.44}$$

$$\psi(z) = \tan^{-1}\left(\frac{z}{z_R}\right) = \tan^{-1}\left[\tilde{q}^{-1}(z)\right]. \tag{3.45}$$

3.5 Modes

This completes a simple framework for how quantum light can be discussed. Using reasonable approximations, we can describe the spatial structure of a laser beam and determine how a short pulse will travel along the beam. We can define a mode, centred about a given carrier frequency, as a slowly-varying envelope in time and space that propagates according to simple equations of motion. If a given quantum state can be prepared in a given temporally-shaped pulse in the laser beam, the linearity of the Maxwell equations ensures that it will propagate in the same manner as the classical vector potential. Operations can be performed on this mode such as delaying it time, mixing it with another mode on a beam-splitter or subjecting it to loss. All of these operations can be treated classically and the result will apply to the quantum state.

Knowing that Gaussian spatial modes are eigenmodes of both free-space propagation and transformations through spherical optics makes it necessary to keep track of only the complex beam parameter to describe the spatial structure of an optical mode. Two laser beams are said to be mode-matched when they share the same spatial mode and frequency, for example, aligning a free-space laser beam into an optical cavity requires adjusting its spatial properties to exactly match the structure of cavity eigenmode as well as tuning the frequency to be on resonance with the cavity mode. Another situation that involves mode-matching occurs when combining two beams on a beam-splitter such that they interfere with each other. For laser pulses, the mode has a defined temporal structure as well as a spatial structure. An optical mode is then defined by the temporal envelope, frequency, and spatial shape of the pulse.

The term mode is used as a convenient shorthand when discussing alignments in optics and especially quantum optics but it refers to a very specific definition: a mode is a unique solution to the wave equation for the vector potential. In the same way that two quantum states are the same if they differ only by a scalar multiplication, two optical modes are the same if they correspond to electric fields that differ only in amplitude. We can apply the same rules that govern quantum states to optical modes. For example an inner product that gives the overlap between two modes can be defined.

CHAPTER 4

Experimental techniques in quantum optics

The aim of this chapter, like the preceding one, is to provide a brief intuitive overview of a handful of topics that are useful in experiments in quantum optics and are used later in this thesis. For a more complete discussion of the topics, I refer the reader to more comprehensive resources on laser optics such as [b6, b7].

4.1 Non-linear optics

The instantaneous polarisation of a dielectric medium is a result of a small displacement of the mean electron position relative to the nuclei. For small displacements, the energy potential for the electrons is approximately harmonic, resulting in a linear response to driving fields. As the driving fields increase in magnitude, however, the electron potentials exhibit anharmonic motion as more of the structure of the potential is revealed. This anharmonic motion results in a non-linear susceptibility to driving electric fields and depends on the structure of the medium as well as the polarisation and frequency of the driving fields.

The complete non-linear response of a medium's polarisation can be approximated by a Taylor expansion :

$$\mathbf{P} = \mathbf{P}_0 + \epsilon_0(\chi^{(1)}(\omega)\mathbf{E}^{(1)}(\omega) + \chi^{(2)}(\omega)\mathbf{E}^{(1)}(\omega)\mathbf{E}^{(2)}(\omega) \tag{4.1a}$$

$$+ \chi^{(3)}(\omega)\mathbf{E}^{(1)}(\omega)\mathbf{E}^{(2)}(\omega)\mathbf{E}^{(3)}(\omega) + \cdots) \tag{4.1b}$$

where ω is the frequency of the light that is illuminating it, ϵ_0 is the vacuum permittivity and $\chi^{(n)}(\omega)$ are rank $n+1$ tensors that relate the orientation of the driving fields to the orientation of the resultant medium's polarisation.

$\chi^{(1)}$ is the linear response of a dielectric medium and describes the (linear) refractive index including birefringence. $\chi^{(2)}$ is the second-order response and can be present only in non-isotropic materials such as optical crystals. This is the most used non-linearity in optics and is used for second harmonic generation, sum frequency generation, difference frequency generation and parametric down conversion. The $\chi^{(3)}$ non-linearity, known as the Kerr effect, is most commonly used in isotropic materials such as glasses or atomic vapours. It is responsible for effects such as third-harmonic generation and four-wave-mixing. $\mathbf{P}_0$ is a polarisation that is intrinsic to the dielectric medium. This is referred to as ferroelectricity and plays a role in non-linear optics because it can be used to alter the electric susceptibility.

Both the $\chi^{(2)}$ and $\chi^{(3)}$ non-linearities are important for this work although they are used in very different contexts. The goal of developing monolithic resonators from $\chi^{(2)}$ non-linear crystals is to perform second harmonic generation and down-conversion. In this case the light is far from any optical resonance in the crystal and interacts weakly with a large number of atoms. The resultant non-linearity acts similarly on a broad range of frequencies.

The non-linear interactions in rubidium vapour that are used for the memory experiments are a $\chi^{(3)}$ effect where the intensity of one field modulates the refractive index for a second field. In this case, however, the interaction takes place very near an atomic resonance and involves relatively few atoms. The effect is strongly frequency dependent and can be controlled to a fine degree by manipulating the internal states of the atoms involved.

4.1.1 Second harmonic generation

When a material that exhibits second-order non-linearity is illuminated by a sufficiently intense field it will result in a polarisation with a component at twice the frequency of the incident light [b7]:

$$\mathbf{P}(2\omega) = \epsilon_0 \chi^{(2)} \mathbf{E}(\omega)\mathbf{E}(\omega). \tag{4.2}$$

The polarisation at 2ω will then produce an optical field at that frequency and we can find the evolution of this second harmonic field from eq. (3.33). Assuming particular polarisations for the fields involved and selecting the appropriate element of $\chi^{(2)}$ for the interaction, the envelope of the second-harmonic field will build according to

$$\frac{\partial A_{2\omega}}{\partial z} = \frac{i k_{2\omega} d_{eff}}{2 n_{2\omega}} A_{\omega}^2(z) e^{i\Delta k z} \tag{4.3}$$

where d_{eff} is the effective non-linear coefficient in m/V and $\Delta k \equiv (2k_\omega - k_{2\omega})$ is the mismatch in the wave-number between the two fundamental waves and the second harmonic due to dispersion in the non-linear medium. The condition that must satisfied for efficient second harmonic generation is that $\Delta k \approx 0$, which is known as the phase-matching condition.

4.1.2 Phase-matching

Phase-matching refers to the matching the wave-number of the fundamental wave with that of the second-harmonic wave in second-harmonic generation (SHG) or parametric down-conversion (PDC) process. This ensures that light generated at any point in the crystal is in phase with the wave that has been produced at previous points along the crystal. More generally, the term can applied to any non-linear conversion process, since momentum conservation demands that the sum of wave-vectors of the initial photons must be equal to the sum of wave-vectors of the generated photon.

Phase-matching is achieved in one of two ways: critical phase-matching and non-critical phase-matching. Critical phase-matching is also known as angle phase-matching. It involves tuning the angle of a birefringent material to match the refractive indices of orthogonally polarised fundamental and second-harmonics. Non-critical phase-matching involves, tuning the temperature of the crystal to achieve the same when the waves of the fundamental and second harmonic lie along the fast and slow axes of the crystal.

Non-critical phase-matching uses the birefringence of an optical crystal to nearly phase-match the fundamental and second harmonic waves for a particular wavelength. The remaining difference in wave-number is eliminated by tuning the temperature of the crystal. This method is restricted to materials that have appropriate birefringence to cancel dispersion at a temperature that can reasonably be achieved in the lab. For a fundamental wavelength around 1064 nm, $\mathrm{LiNbO_3}$ is such a material.

A third option is to not phase-match the interaction at all, but to use periodic poling of the crystal. The technique, known as quasi-phase-matching, is to alternate the direction of the crystal's optical axis such that as the phase mismatch between the fundamental and second harmonic passes $\pi/2$ the sign of the non-linear interaction is inverted.

4.2 ABCD matrices

Ray propagation matrices, also known as ABCD matrices, are a tool for calculating the propagation of an optical ray through a system of optical elements under the assumptions that the ray always forms only a small angle relative to the optical axis and that the ray always propagates within a plane. At any given point along the optical path, the ray trajectory can be described by two quantities, the distance of the ray from the optical axis, $r(z)$, and the angle of the ray relative to the optical axis $r'(z)$. We can list these quantities into a vector, and describe the a transformation through an optical system using a matrix:

$$\mathbf{r_2} \equiv \begin{pmatrix} r_2 \\ r_2' \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} r_1 \\ r_1' \end{pmatrix} \equiv \mathbf{M} \cdot \mathbf{r_1}. \tag{4.4}$$

This formalism doesn't apply to arbitrary optical elements, but it does work for those commonly used in optical systems: mirrors, lenses, and free-space propagation. By taking

$$r' \equiv n(z)\frac{dr(z)}{dz} \tag{4.5}$$

the matrices also apply to propagation through boundaries between materials with different refractive indices [b6]. The ABCD matrices associated with the basic elements needed to analyse optical resonators are shown in table 4.1. The ray matrices for each element describe a how a ray is transformed from the input of the element to the output. Complex optical systems can therefore be described by the product of the associated elements.

When designing optical resonators, we are concerned with the properties of Hermite-Gaussian beams rather than discrete rays. These beams are eigenfunctions of free-space propagation and also of the transformations associated with simple paraxial optics. As a result, they are also eigenfunctions of systems of paraxial optics and it is because of this that they are of interest. Laser resonators are built from simple paraxial optical and thus the stable modes that are produced are eigenfunctions of those elements. Because of this, we can describe the propagation of Hermite-Gauss modes through a system of paraxial optics entirely through the complex beam parameter, $\tilde{q}$, that describes the size and wave-front curvature of the mode.

The transformation of complex beam parameter through a system of paraxial optical elements can be found through the ABCD matrix for that system. The derivation can be found in [b6], the result of which is that $\tilde{q}$ transforms according to

$$\frac{\tilde{q}_2}{n_2} = \frac{A(\tilde{q}_1/n) + B}{C(\tilde{q}_1/n1) + D}. \tag{4.6}$$

(a) Free-space propagation	$M_{fs}(L) = \begin{pmatrix} 1 & L/n \\ 0 & 1 \end{pmatrix}$
(b) This lens, focal length f	$M_l(f) = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}$
(c) Curved mirror, arbitrary incidence	$M_{cm}(R_e(\theta)) = \begin{pmatrix} 1 & 0 \\ -2/R_e & 1 \end{pmatrix}$
(d) Curved dielectric interface, tangential plane	$M_{ci,t}(\theta_i, \Delta n_e) = \begin{pmatrix} \frac{\cos\theta_2}{\cos\theta_1} & 0 \\ \Delta n_e/R & \frac{\cos\theta_1}{\cos\theta_2} \end{pmatrix}$
(e) Curved dielectric interface, sagittal plane	$M_{ci,s}(\theta_i, \Delta n_e) = \begin{pmatrix} 1 & 0 \\ \Delta n_e/R & 1 \end{pmatrix}$

Table 4.1: ABCD matrices for simple optical elements. For (c), the effective radius of curvature is $R_e(\theta) = R\cos\theta$ for the tangential plane and $R_e(\theta) = R/\cos\theta$ for the sagittal plane.

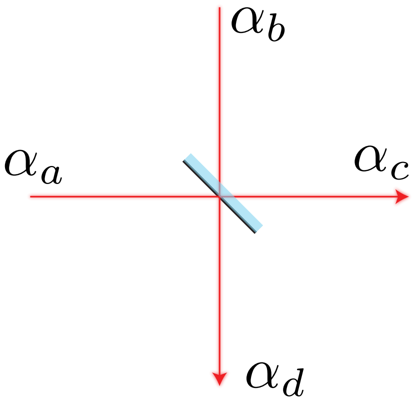

Figure 4.1: A simple conceptual model of a beam-splitter. The side of the beam-splitter with the π phase shift for reflected light is indicated with a black line.

We can apply this to determine how a Hermite-Gaussian beam will propagate through complex optical systems. The eigenmode of an optical resonator can also be found by solving

$$\tilde{q} = \frac{A\tilde{q} + B}{C\tilde{q} + D}, \tag{4.7}$$

where the ABCD matrix describes all the optical elements and free-space propagation for one round trip through the cavity.

4.3 Beam-splitters

Two spatial modes can be combined on a beam-splitter simply by ensuring that they have the same size and divergence at the location of the beam-splitter interface. This ensures that they overlap completely at both outputs of the beam-splitter. The combination of the quantum modes can then be easily described.

The beam-splitter performs a simple transformation on two optical modes, combining them into two new superpositions of the input modes. Assuming that two input modes with classical amplitudes, α_a and α_b, are mode-matched, the output modes, α_c and α_d,

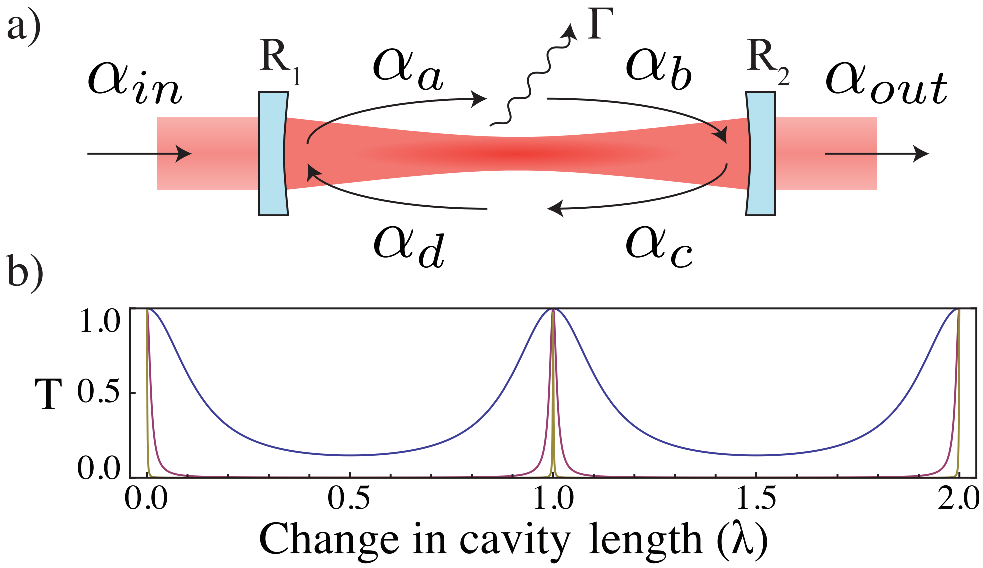

Figure 4.2: a) A simple conceptual model of an optical cavity. b) Transmission spectrum for a lossless ($\Gamma = 0$) optical cavity with $R = 0.47$ (blue), $R = 0.94$ (red), and $R = 0.99$ (yellow).

are

$$\begin{pmatrix} \alpha_c \\ \alpha_d \end{pmatrix} = \begin{pmatrix} t_{ac} & r_{bc} \\ r_{ad} & t_{bd} \end{pmatrix} \begin{pmatrix} \alpha_a \\ \alpha_b \end{pmatrix} \tag{4.8}$$

where t is the transmittance and r is the reflectance. Figure 4.1 shows an illustration of the model used. Energy conservation requires the relations

$$|r_{ac}|^2 + |t_{ad}|^2 = |r_{bd}|^2 + |t_{bc}|^2 = 1 \tag{4.9}$$

and

$$r_{ac}t_{bc}^* + t_{ad} + r_{bd}^* = 0. \tag{4.10}$$

This is often simplified by selecting a global phase shift for the outputs such that

$$\begin{pmatrix} \alpha_c \\ \alpha_d \end{pmatrix} = \begin{pmatrix} t & r \\ -r & t \end{pmatrix} \begin{pmatrix} \alpha_a \\ \alpha_b \end{pmatrix}. \tag{4.11}$$

4.4 Optical cavities

The development of compact optical cavities is the focus of the second part of this thesis. Cavities are used for frequency filtering in the memory experiments and frequency calibration in the resonator experiments. Beyond these explicit applications in the thesis, they are responsible for producing coherent light from all of the lasers used.

A simple model for an optical cavity can be derived from the model shown in fig. 4.2

(a). The equations for the amplitudes of the fields shown are

$$\alpha_a = -r\alpha_d + t\alpha_{in}, \qquad \alpha_b = e^{ik\Delta L}\alpha_a, \tag{4.12a}$$

$$\alpha_c = r\alpha_b, \qquad \alpha_d = (1-\Gamma)e^{ik\Delta L}\alpha_c, \tag{4.12b}$$

$$\alpha_{out} = t\alpha_b, \qquad t = \sqrt{1-r^2}, \tag{4.12c}$$

with the solution

$$\frac{\alpha_{out}}{\alpha_{in}} = \frac{e^{ik\Delta L}(R-1)}{1+e^{ik\Delta L}R(e^{ik\Delta L}-\Gamma)} \tag{4.13}$$

where $R = r^2$.

The intensity transmission of a lossless, $\Gamma \to 0$, cavity is

$$T = |\frac{\alpha_{out}}{\alpha_{in}}|^2 = \frac{(1-R)^2}{1+R^2+2R\cos(2\Delta)} \tag{4.14a}$$

$$= \frac{1}{1+F\sin^2\Delta/2}, \text{ where } F = \frac{4R}{(1-R)^2} \tag{4.14b}$$

and is shown in fig. 4.2 (b). The finesse $\mathcal{F}$ is related to the coefficient of finesse by

$$\mathcal{F} = \frac{\pi}{2\text{arcsin}(1/\sqrt{F})} \approx \frac{\pi\sqrt{R}}{1-R}. \tag{4.15}$$

The finesse is a measure of how many times a photon can circulate around the cavity prior to being lost through the mirrors. The cavity linewidth can be expressed through the finesse as well as the free spectral range by

$$\delta = \frac{\text{FSR}}{\mathcal{F}}. \tag{4.16}$$

The free spectral range $FSR = c/L$, where L total path-length required for a photon to complete one round-trip of the cavity.

Another quantity that is used to describe optical cavities is the quality factor, or $\mathcal{Q}$-factor. This is defined as

$$\mathcal{Q} = \frac{\nu}{\delta} \tag{4.17}$$

where ν is the optical frequency. The quality factor is proportional to the mean number of oscillations of the electric field before a photon in the cavity decays. This measure is useful when comparing cavities that have losses that are predominantly due to absorption along the entire path length rather than only at mirrors or scattering centres.

The transmission through the cavity shows sharp resonances that occur when the optical path length is an integer number of wavelengths of the input laser and the width of these resonances is determined by the reflectivity of the mirrors. By using high reflectivity mirrors, the resonances can be used for frequency filtering because the cavity will reject optical frequencies that are not resonant. On resonance, the amplitude of the optical fields within the cavity is increased by a factor of $\mathcal{F}$, a property that has significant usefulness for non-linear optics.

4.5 Locking

Obtaining a high transmission efficiency through an optical cavity requires mode-matching the input beam to the cavity mode. Spatially, this is done by adjusting the size and divergence of the input. The frequency of the input laser must also be matched to a cavity resonance, a task that is usually performed by actively stabilising the laser frequency to the cavity. The most common technique to do this is the Pound-Drever-Hall (PDH) locking technique. An excellent review of the technique can be found in [a8].

The challenge of locking a laser beam to a cavity resonance lies in the fact that the reflection spectrum $R(\Delta) = 1 - T(\Delta)$ is symmetric about resonance so a decrease in the reflection from resonance gives no information about whether the input frequency should be increased or decreased to correct the change. The insight that leads to PDH locking is that the derivative of $R(\Delta)$ is anti-symmetric about resonance. The derivative of $R(\Delta)$ can be accessed by rapidly varying the detuning Δ. Measuring the magnitude of change in $R(\Delta)$ when Δ increases or decreases by known amount provides a signal that is proportional to $R'(\Delta)$, providing an error signal for the input frequency.

If the modulation of the input frequency is sufficiently fast, the error signal can be extracted by mixing the detection signal from the cavity reflection with the signal used to modulate the input frequency and sending the result through a low-pass filter. The frequency of a laser can be stabilised to an optical cavity resonance or an atomic resonance by taking the error signal from the PDH technique and feeding it back to the laser through a feedback controller. Equivalently, a cavity can be stabilised to a laser using the same technique.

4.6 Feedback control

The most common type of feedback control loop used in the experiments is the PID controller, referring to proportional, integral, and derivative components of the error signal. The proportional component of the controller provides feedback that is proportional to the error signal, the integral component provides feedback that is the integrated total of the error signal in the past, and the derivative component gives feedback that is proportional to the rate of change of the error signal in time. The feedback $f(t)$ is given by the sum of these three components

$$f(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \partial_t e(t) \tag{4.18}$$

where $e(t)$ is the error signal and $K_{\{p,i,d\}}$ are the gains of the three components. These gains depend on the dynamics of the system that is being controlled and are chosen to provide the fastest settling time.

PID controllers are used in the experiments for locking optical filtering cavities to laser fields, locking laser frequencies to atomic resonances, and temperature control of optical elements. The implementation of the controllers can be done with analogue electronics, but digitising the error signal and implementing the controller in software provides advantages in terms of flexibility. The cavity locks used in the hot-atom memory experiment are custom analogue controllers that use only proportional and integral feedback, the locks in the cold-atom memory experiment are entirely digital and the temperature control for the resonator experiment uses a commercial PID controller that provides all three gain components and the option to tune them automatically by measuring the system response.

4.6.1 Digital control

A majority of the feedback controllers used for experiments in this thesis are implemented in software using the LabView programming environment. For this approach, the error signal is digitised, passed to an FPGA (field-programmable gate array) which passes the signal through the control algorithms and the output is again converted to an analogue feedback signal and fed to the system. The FPGAs used are integrated into cards that contain a set of analogue-to-digital converters (ADCs) and digital-to-analogue converters (DACs) that can be used to pass the error signal into the FPGA and the feedback signal out to the experiment. Each FPGA has the capacity to run up to eight locks simultaneously and the general PID code can be used with only moderate tailoring to different locks.

The digital locking system used in our experiments was initially developed by Ben Sparkes and a description of it can be found in [a9, c10]. The digital control system was further developed by Seiji Armstrong and a complete guide for implementing feedback controllers in LabView can be found in his thesis [c11] along with sample code.

4.7 Electro-optics

One of the important techniques in quantum optics is modulation of the optical fields. This is required for the frequency modulation used in locking techniques as well as for general control of the frequency phase and amplitude of light in atom optics experiments. The two types of optical modulator that we use are the acousto-optic modulator and the electro-optic modulator.

4.7.1 Acousto-optic modulators

Experiments in optical memory require pulsing and fine-frequency tuning. Both of these tasks are accomplished by acousto-optic modulators (AOMs) which produce a tunable, frequency shifted output typically in a 10 MHz bandwidth around 80, 110, or 200 MHz. These devices use high-frequency sound waves in optical crystals to diffract laser beams. The acoustic compression and rarefactions in the crystal produce refractive index variations that act as a diffraction grating propagating transversely to the laser beam. Because the grating is moving, the diffracted light is frequency shifted. The AOMs can be turned on an off quickly to generate pulses and can be used to shape pulses and control the frequency and phase of the diffracted beam precisely.

The output beam emerges at a diffraction angle proportional to the frequency shift. To avoid alignment errors while adjusting the frequency, the modulators can be double-passed: the output is collimated and retro-reflected back into the AOM resulting in twice the frequency shift with change in alignment for varying frequencies. A polarising beam-splitter and $\lambda/4$ wave-plate are used to separate the input light and backward-propagating, frequency shifted light.

4.7.2 Electro-optic modulators

AOMs are limited to a narrow bandwidth by the ultra-sonic transducer that drives the acoustic waves. For more broadband modulation, electro-optic modulators (EOMs) can be used. These use the electro-optic effect to modulate the refractive index of a crystal in response to an applied electric field. Bulk crystal can produce modulations up to a few hundred MHz, limited by the capacitance of the crystal and electrodes. High-frequency EOMs are made from crystal wave-guides that use thinner crystal to provide a higher electric field for a given voltage. This, combined with a lower capacitance due to a small area allows fibre-coupled wave-guide EOMs to provide modulations up to 10 GHz.

Bibliography - Introduction

Articles

a8. Black, E. D. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics* **69,** 79–87 (2001).

a9. Sparkes, B. *et al.* A scalable, self-analyzing digital locking system for use on quantum optics experiments. *Review of Scientific Instruments* **82,** 075113 (2011).

Books

b1. Walls, D. F. & Milburn, G. J. *Quantum Optics* ISBN: 0-387-57179-5 (Springer-Verlag, 1994).

b2. Leonhardt, U. *Essential Quantum Optics: From Quantum Measurements to Black Holes* ISBN: 9780521869782 (Cambridge University Press, 2010).

b3. Marchildon, L. *Quantum Mechanics* ISBN: 3-540-43342-2 (Springer-Verlag Berlin Heidelberg, 2002).

b4. Lambropoulos, P. & Petrosyan, D. *Fundamentals of Quantum Optics and Quantum Information* ISBN: 978-3-540-34571-8 (Springer Berlin Heidelberg NewYork, 2007).

b5. Jackson, J. D. *Classical Electrodynamics* Third. ISBN: 0-471-30932-X (John Wiley & Sons, 1998).

b6. Siegman, A. E. *Lasers* ISBN: 0-935702-11-3 (University Science Books, 1986).

b7. He, G. S. & Liu, S. H. *Physics of nonlinear optics* ISBN: 9810233191 (World Scientific Publishing, 1999).

Theses

c10. Sparkes, B. M. *Storage and Manipulation of Optical Information Using Gradient Echo Memory in Warm Vapours and Cold Ensembles* PhD thesis (The Australian National Univeristy, 2013). <`http://photonics.anu.edu.au/theses/2013Sparkes.pdf`>.

c11. Armstrong, S. *Experiments in Quantum Optics: Scalable Entangled States and Qauntum Computation with Cluster States* PhD thesis (The Australian National Univeristy, 2014).

I

Quantum memories

CHAPTER 5

Overview of quantum memory and applications

In this part of the thesis I will present our results regarding quantum memories in rubidium vapours. The goal of the work is to develop high-efficiency optical memories that can store and recall quantum states with a high fidelity. The results show that we can achieve these goals with memories based on either electromagnetically induced transparency or gradient echo memory. I then proceed to present a number of protocols to extend the functionality of the gradient echo memory scheme to tasks in quantum information processing.

5.1 Overview of quantum memory and applications

Optical quantum memories are an integral component of secure quantum communication and quantum computation [a1, a2]. Many implementations of memories are being developed and much experimental progress has been made, with demonstrations of efficiencies approaching 90% [a3–a5], storage times of over seconds [a6–a8], bandwidths of hundreds of MHz [a9, a10] and large multi-modal capacities [a3, a10, a11].

There are a wide variety of techniques for implementing optical quantum memories. The simplest are passive linear optical systems such as fibre delay lines or optical cavities that delay an input pulse for a fixed duration. Such devices are a mainstay of quantum information processing and are used to synchronise photons in optical networks or to provide time for electro-optic control elements to perform conditional operations. More sophisticated memories generally have the ability to control the duration of the delay before the optical pulse exits the memory. A simple switch on a loop of fibre or an electro-optic element in an optical cavity may suffice for short-duration memories but longer duration storage times require writing the information contained in the optical pulse into some stationary medium.

The majority of proposed and demonstrated memories solve this problem by writing the optical information into the internal state of an atom, a collection of atoms, or an atom-like quantum system. For a single atom, a large coupling between the atom and a single photon is needed, usually by confining both to a small region to improve the coupling rate. An alternative approach is to use an ensemble of atoms that contains a large enough number of atoms to completely absorb a passing photon even if each atom interacts only weakly with each photon. An atomic ensemble-based optical memory therefore relies on the absorption and coherent re-emission of an optical field by the entire collection of atoms.

The write phase of the memory is straightforward, shining the light into an absorbing

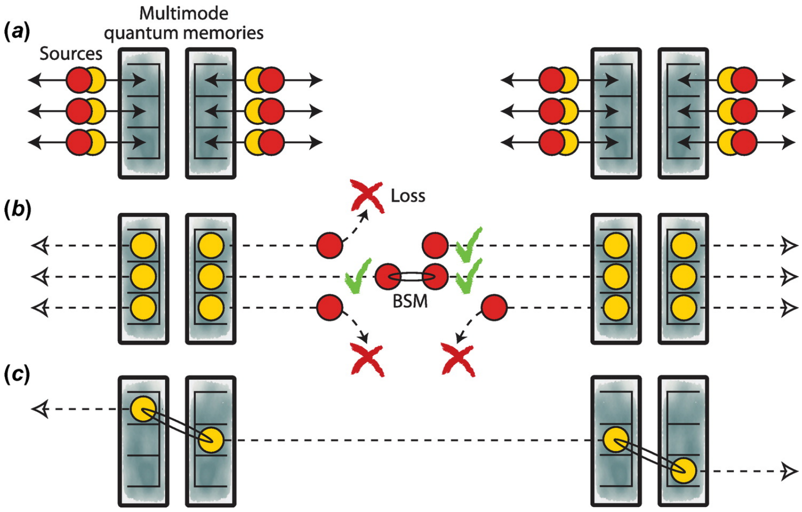

Figure 5.1: A conceptual schematic of a quantum repeater. Each node of the repeater contains two sets of quantum memories that each store one of the outputs of a source of entangled photon pairs. The other side of each entangled photon pair is sent to an intermediate node for measurement. When a successful joint measurement is made between two photons, the contents of the memories at the nodes that sent them become entangled. The nodes can be linked by recalling the contents of memories that contain successfully entangled photons at each node and performing a joint measurement to swap the entanglement to the adjoining nodes. The figure is reproduced from [a2]

sample will be convert the light into just such a coherence. The read phase, however, requires some effort. Simply waiting for the sample to re-emit will produce incoherent light in a random direction that contains little of the information that was encoded in the optical field. Coherent re-emission requires some manipulation of the sample such that the information contained in the atomic coherence is converted back into an optical field before spontaneous emission erodes the state. There is a wide variety of proposals to engineer the dynamics of absorption and re-emission from an atomic ensemble most of which fall into a few general techniques. These include slow-light [a12–a14], photon echoes [a15, a16], and time-reversal symmetries [a17–a19], although individual memory protocols may be combine these techniques.

5.2 The motivation for developing a quantum memory

There is a wide variety of physical systems that are being explored as candidates for performing quantum computation. However, while optics is only one among many possible architectures for quantum computing, it is the only viable candidate for quantum communication over large distances due to the fact that photons travel quickly and interact weakly with the environment. For this reason, developing a network for quantum information is relevant for linking together quantum computers of any type. The difficulty of building a network on a large scale is the loss of photons due to absorption or scattering. This loss, unlike loss in a classical information network cannot be corrected by simple amplification techniques because these add intrinsic noise to the quantum state.

The solution that has been proposed is to teleport information directly from the sender to the receiver, thereby avoiding loss in the transmission lines [a2, a20, a21]. In this scheme, a series of short-range quantum repeater nodes is used to generate entanglement as a resource to be used to teleport the quantum states of interest. Each repeater node contains two optical quantum memories and a source of heralded entangled photon pairs. The node will aim to store one photon of a locally generated entangled pair and send the other to the next node in the chain. That next node will do the same and store any photons arriving from the previous node. When a node has a photon stored from its own entangled source and one from the previous node in the chain, it can recall both photons and perform a joint measurement on them. This measurement will swap the entanglement to form an entangled stated between the previous node and the next one. The idea is illustrated in fig. 5.1. Loss in the fibre links will decrease the success rate of the entangling operations but for linking entanglement over long distances the use of memories can, in principle, yield significant improvements over direct transmission [a21].

A further incentive for developing quantum memories is to develop resources for linear optics quantum computing. This scheme for quantum computation relies on measurement-based quantum gates that can yield large non-linear interactions between single photons in a probabilistic manner [a22]. The probabilistic nature of the gate operation requires a large number of single photons to be input into the linear gates [a23]. Because the generation of single photons themselves is probabilistic, quantum memories can be used to store generated single photons until a sufficient number are available for a computation and release them into an optical network simultaneously [a2]. Additionally, memories could be used to store the results of successful gate operations for use later in an optical circuit, dramatically improving the scalability of a linear optical quantum computer [a24].

5.3 Slow light

Slow-light is perhaps the most conceptually straightforward of the memory protocols. The idea is that if the dispersion of a material can be engineered to exhibit spectrally narrow features, the group velocity of optical pulses centred about those features will be reduced and the pulse will be delayed. If the delay can be controlled then the slow-light system can be used as a buffer to synchronise the arrival times of optical pulses. Electromagnetically induced transparency provides a means to engineer a narrow transparency window in an otherwise absorptive atomic transition line. This provides the spectral feature necessary to reduce the group velocity of an incident pulse.

Electromagnetically induced transparency is a quantum effect that takes place in a three-level atomic structure that is driven by a strong optical field [a12, a13]. The electronic level structure involves a ground state, and excited state, and an auxiliary meta-stable state. Both the ground state and the meta-stable state have optical transitions to the excited state. If a bright field drives the transition between the meta-stable state and the excited state, a narrow transparency feature will open in the transition line between the ground state and excited state.

Interest in slow-light as a means to achieve a viable optical quantum memory increased when it was shown that optical information contained in an EIT slow-light pulse can be transferred to the atoms by turning off the bright driving field. The process induces signal photons to coherently scatter off of the excited state and transfer information into the coherence between the ground and meta-stable states. The fact that the scattering process is coherent, that is to say, the excitation induced in the atoms has a phase directly related to that of the scattered photon, means that the process can be reversed. This reversal is done by restoring the driving field to its original intensity, which induces a scattering of driving field photons back into the original signal field.

A substantial amount of literature has been written on the EIT effect and its application to quantum memory. Since EIT was first described experimentally in a strontium vapour in 1991 [a25], it has been investigated in a wide variety of atomic systems for a wide variety of applications. It is of fundamental interest for its ability to slow light by up to seven orders of magnitude [a26, a27]. In turn, the increased interaction times afforded by slow light allows enhanced non-linear optical interactions [a28–a30]. EIT can also be used to stop and store light in an atomic spin-wave leading to applications as a quantum memory [a1, a14, a31–a33]. For EIT to be used as a quantum memory the process must be both efficient and noiseless [a34]. Considerable work has been done to improve the efficiency, primarily by increasing the optical depth of the slow-light medium and recent results have demonstrated up to 69% efficiency for storage and forward recall of light [a35]. In principle, the achievable storage efficiency using EIT can approach 100% [a19] as the optical depth is made sufficiently large. In this regime, the delay of an optical pulse travelling through the EIT medium can be larger than the duration of the pulse [a36], such that it is contained entirely within the medium. Chapter 8 contains a more detailed description of EIT slow-light and memory.

5.4 Controlled reversible inhomogeneous broadening

Another common approach to optical quantum memory is to exploit time-reversal symmetries in the equations that govern the light-atom interaction. The principle of controlled reversible inhomogeneous broadening is that if the decay process responsible for dissipating the atomic excitation associated with the absorption of a photon can be reversed, it is possible to recover the photon [a15, a37, a38]. Unless the dissipative processes are carefully engineered, however, they are not usually reversible. Spontaneous emission or relaxation of the atomic coherences due to external noise sources results in irrecoverable losses. The usual approach is to try to suppress these mechanisms or operate on time-scales shorter than the incoherent decay rates and then intentionally apply an engineered de-phasing mechanism that can be reversed, hence the name controlled reversible inhomogeneous broadening (CRIB).

The most common approach is to reduce the inhomogeneous broadening of a large ensemble of atoms, either by eliminating mechanisms that introduce broadening or by removing all atoms that fall outside of a natural linewidth from the central frequency of the protocol. Eliminating broadening mechanisms may involve eliminating impurities from solid-state systems because they introduce different local electric and magnetic fields for the atoms of interest. Another option is to select transitions that are insensitive to the broadening mechanism, for example using two-photon transitions in warm vapours that are insensitive to Doppler broadening. Removing unwanted atoms from an ensemble is usually accomplished with spectral hole burning, which uses a pump laser to prepare atoms that are at incorrect detuning into electronic states that do not absorb the optical frequency of interest.

Once the ensemble is prepared such that only a narrow linewidth remains, the line is then broadened by the application of a spatially inhomogeneous external electric or magnetic field. Optical pulses with a bandwidth that matches that of the broadened linewidth are absorbed fully into the atomic ensemble and the induced coherence rapidly de-phases due to the broadening of the linewidth. Unlike a decay process such as spontaneous emission, however, the magnitude of the coherence at any given location in the ensemble does not decay, only the collective excitation of the entire ensemble is suppressed due to different local phases accumulated at each location. If the relative sign of the detunings of the atoms can be reversed the ensemble will re-phase and the collective excitation will be recovered.

Two methods are used to accomplish the re-phasing. One option is to apply a π-pulse to the ensemble to invert the population. The broadening will now affect the local phase of the atomic coherence with an opposite sign and the ensemble will re-phase. Another method is to invert the sign of the externally applied broadening field, which has the same effect on the phase of the coherence but without the need to invert the population.

An intuitive picture for the re-phasing process is to imagine a group of runners lined up alongside each other. At the start of the race, the runners are all at the start line but depart with varying speeds depending on their abilities. After some time the group will become spread out with the faster runners farther from the start line. If at some time a whistle is blown and all the runners reverse direction, they will run back to the start line each at their own speed. The faster runners, however, will cover the longer distance back to the start in the same time that the slow runners take to to cover the shorter distance. All the runners therefore return to the start line simultaneously. This is analogous to the de-phasing of atoms in an ensemble that has an applied inhomogeneous broadening; each atom will accumulate phase at a rate different to other atoms in the ensemble according to its particular detuning. When the sign of the broadening is reversed, the atoms will begin to re-phase and recover the initial coherence.

Once the ensemble re-phases, the collective excitation will generate a polarisation that is identical to that when the photon was initially absorbed, recreating the photon itself. The challenge then, is to avoid re-absorption of the photon. This can be done by reversing the direction of propagation. In this case, the backwards recall of the photon is the exact time-reversal of the absorption process and this symmetry ensures efficient recall. Achieving this backwards recall can be accomplished by the application of multiple π-pulses with different propagation direction. This cycles the coherence through auxiliary energy levels but prepares the coherence with a backward-propagation wave-vector.

5.5 Gradient echo memory

A majority of the work that I have done on quantum memories has focused on gradient echo memory (GEM). GEM is a particular implementation of controlled reversible inhomogeneous broadening in which the applied inhomogeneous broadening takes the form of a linear gradient of detunings in the ensemble, applied along the direction of propagation [a4, a18, a39, a40]. Each frequency component of the incoming pulse is therefore encoded in a particular slice of atoms along the length of the ensemble, effectively spatially encoding the Fourier transform of the input pulse in the ensemble. The broadening causes the atomic excitation to de-phase, and the excitation is frozen into the atoms. If the gradient is reversed at a later time, the ensemble will begin to re-phase and re-emit the original pulse. This implementation of CRIB has the advantage that re-absorption is suppressed. The reversal of the gradient has a similar effect to backwards retrieval in standard CRIB schemes; it creates a time-reversal symmetry for recall in the forward direction that ideally recreates the input pulse at the exit of the memory. This symmetry also has the effect of time-reversing the pulse, so that if multiple pulses are stored, the last one in will be the first one out.

The spatial encoding of frequencies provides GEM with a further advantage over standard CRIB schemes which is that spectral processing can be performed using the memory [a11, a41]. Basic operations such as temporally compressing or expanding the recalled pulse can be performed by altering the magnitude of the gradient used for the recall stage. More complex operations can be performed by modifying the gradient into more complex functions, allowing individual Fourier components of the input pulse to be stored or recalled independently.

5.6 Atomic frequency comb memories

The memories previously discussed ideally rely on an ensemble of atoms with a very narrow initial linewidth because any broadening of the linewidth limits the storage lifetime of the memory. In some systems, such as atomic vapours, relatively narrow line-widths are easily accessible for transitions between spin states. For other systems however, the ensemble has a very large intrinsic inhomogeneous broadening. An important example is solid-state systems. These use ensembles of ions that are doped into a crystal substrate and are subject to large inhomogeneous broadening due to varying local electric fields throughout the crystal [a42]. These local fields are due to uneven distribution of isotopes, impurities or defects. CRIB memories in these systems require the linewidth of the ensemble to be narrowed by spectral hole-burning, using a bright pump laser to drive all atoms that have transitions at unwanted frequencies into dark states that do not interact with the light of interest. The downside to this technique is that a large portion of atoms in the ensemble are unused, reducing the effective optical depth, and thus the performance of the memory. A solution to this is the atomic frequency comb memory.

In this memory, spectral hole burning is not used to reduce the ensemble to a single narrow spectral feature, but rather to a comb of narrow features [a43, a44]. In this way, far fewer atoms are rejected from the ensemble. The comb of absorption lines gives a periodic spectral response in the frequency domain, which corresponds to a periodic impulse response in the time domain. This periodic response, combined with the normal dispersion in the vicinity of the absorption features, results in a series of echo pulses when an input pulse is incident on the ensemble. The large inhomogeneous linewidths available provide exceptionally high bandwidths for the associated quantum memories [a9, a10]. The protocol has been demonstrated to work in rare-earth ion doped crystals [a9, a45] at and erbium-doped optical fibres operating at telecom wavelengths [a10]. The AFC memory provides only a fixed time delay based on the spacing of the comb teeth of the absorption features, however, the excitation can be transferred to a long-lived spin-state using π-pulses for a true memory [a45].

5.7 Raman memory

Off-resonant Raman memory is related to EIT in that it uses stimulated Raman scattering to couple a weak signal pulse to a spin transition. In the case of an off-resonant Raman memory, the bright control field is not used on resonance to create a transparency window, but rather off-resonance to create a Raman absorption line. The spectral broadening required to absorb a short pulse is provided by using a spectrally broad control field, accomplished by matching the temporal shape of the control field to that of the signal pulse [a46, a47]. If a second control pulse is sent into the ensemble in the reverse direction, the input pulse is retrieved, also in the backward direction, due to the same type of time reversal symmetry used in GEM and CRIB. If the second control pulse is sent in the forward direction, the protocol works, but with reduced efficiency due to re-absorption. The protocol has been shown to be efficient, broadband [a47] and flexible [a48].

It is interesting to note a similarity between the atomic frequency comb (AFC) memory and the forward-retrieved Raman memory. In the Raman memory, storage is usually performed using a series of control pulses which, in the frequency domain, produce a comb of Raman absorption lines. Much as with the AFC memory, an input pulse is retrieved as a series of output pulses as subsequent control pulses each retrieve a portion of the remaining spin coherence left in the memory. In the Raman memory, however, the absorption lines are coherently related to each other, resulting in an additional condition for storage: the input pulse must arrive at the memory synchronously with a control pulse.

5.8 Context of this work

Development of optical quantum memories has been a widely researched topic in the last decade. While there are numerous figures of merit for quantum memories, efficiency is certainly one of the most important aspects and the cold-atom gradient echo memory described in this thesis currently holds the record for the highest efficiency yet achieved. The two main goals of my work on the topic have been focused on demonstrating the quantum performance of the memory and extending the usefulness of the memory beyond passive storage. Work towards the first goal included the homodyne characterisation [a49, c50] and understanding the four-wave-mixing process in high optical depth ensembles by investigating electromagnetically-induced transparency [a51] in cold atomic ensembles [a5, c52]. Work on the second goal includes the beam-splitting demonstration using the hot-atom memory [a53], the proposals to extend those demonstrations to large number of modes [a54], and the work on storing dual-rail frequency states in the memory for high-fidelity storage.

CHAPTER 6

The theory of light-atom interaction

Memory protocols require fine control over the frequency of an optical transition. This, along with the requirement that the coherence time of the transition be longer than the storage time, implies that the natural inhomogeneous broadening of the transition must be small relative to the time-scales of interest. While this can be satisfied in a few novel materials, it may be simpler in practice to use a spin-coherence between two hyperfine-split ground states for storage. These states are easily found in the alkaline metals, which can be optically addressed if they are prepared as a vapour. The transitions between spin-states are not directly optically addressable so an auxiliary bright optical field is used in order to address them via a two-photon scattering process.

The two-photon Raman scattering process is controlled via a bright, classical coupling field which drives the dipole transition between the meta-stable state and the excited state. A weak, potentially quantum, optical field drives the dipole transition between the ground state and the excited state. Together, the two fields complete a Raman transition between the ground and meta-stable state.

I will provide only a brief derivation of the aspects of light-atom interaction that are important for this thesis. A more complete treatment can be found in [b55]. An easy starting point for exploring the dynamics of light-atom interaction is to start with the simplest system, a hydrogen-like two-level atom driven by a plane, monochromatic optical field. The approach is general, however, and can be extended to higher-dimensional systems.

6.1 The Jaynes-Cummings model

The Jaynes-Cummings model is used to study the interaction between quantum light fields and atomic systems. The approach is to solve the Shrödinger equation for the Hamiltonian associated with the entire light-atom system. This system Hamiltonian,

$$\hat{H} = \hat{H}_{atom} + \hat{H}_{field} + \hat{H}_{int}, \tag{6.1}$$

consists of the sum of the atomic and field Hamiltonians plus an interaction Hamiltonian that quantifies the energy associated with the atom-field coupling. We consider an atom with a ground state $|g\rangle$ and an excited state $|e\rangle$, separated by an energy of $\hbar\omega_0$ where ω_0 is the atomic transition frequency. The Hamiltonian associated with the atom can be written simply as

$$\hat{H}_{atom} = \hbar\omega_0\hat{\sigma}_{ee} \tag{6.2}$$

by defining the energy of the ground state to be zero. Similarly, the Hamiltonian for the optical field is given by

$$\hat{H}_{field} = \hbar\omega\hat{a}^\dagger\hat{a} \tag{6.3}$$

assuming a monochromatic field of frequency ω and setting the zero-point energy to zero for convenience. The interaction Hamiltonian $\hat{H}_{int}$ is derived by considering the energy of the atomic dipole moment $\boldsymbol{d}$ interacting with the optical electric field $\boldsymbol{E}$, classically given by

$$U = -\boldsymbol{d}\cdot\boldsymbol{E}. \tag{6.4}$$

For the quantum interaction, we define a dipole operator $\hat{\boldsymbol{d}} = -e\hat{\boldsymbol{x}}$ that is proportional to the electronic position operator. The diagonal elements of the dipole $\langle g|\hat{\boldsymbol{x}}|g\rangle$ and $\langle e|\hat{\boldsymbol{x}}|e\rangle$ vanish because $|e\rangle$ and $|g\rangle$ are eigenstates of the parity operator and $\hat{\boldsymbol{x}}$ is anti-symmetric. The reason that $|e\rangle$ and $|g\rangle$ have a defined parity (i.e. they are either symmetric or anti-symmetric) is that they are eigenstates of the spherically symmetric potential of the atom [b56]. Intuitively, the symmetry implies that an unperturbed atomic orbital has no mean displacement relative to the nucleus.

The off-diagonal elements of the dipole operator, $\langle e|\hat{\boldsymbol{d}}|g\rangle \equiv \boldsymbol{d}_{eg} = \boldsymbol{d}^*_{ge}$, describe the transition dipole moment for the $|e\rangle \to |g\rangle$ transition. Phenomenologically, this can be interpreted as the mean perturbation of the electron wave-function associated with the change of state and is related to the overlap of the electronic orbitals of $|e\rangle$ and $|g\rangle$.

The Hamiltonian associated with the interaction between the atomic dipole moment and an electric field is

$$\hat{H}_{int} = -\hat{\boldsymbol{d}}\cdot\hat{\boldsymbol{E}}, \tag{6.5a}$$

$$= \hbar g\left(\hat{\sigma}_+ + \hat{\sigma}_-\right)\otimes\left(\hat{a}^\dagger + \hat{a}\right) \tag{6.5b}$$

where $\hat{\sigma}_+$ and $\hat{\sigma}_-$ are the raising and lowering operators $|e\rangle\langle g|$ and $|g\rangle\langle e|$ respectively, and $(\hat{a}^\dagger + \hat{a})$ is the real quadrature of the optical field, eq. (3.23a). The coupling rate g is defined by

$$g \equiv \frac{1}{2}\sqrt{\frac{\omega}{2\hbar\epsilon_0 V}}\left(\boldsymbol{d}_{eg}\cdot\boldsymbol{\epsilon}\right) \tag{6.6}$$

where $\boldsymbol{\epsilon}$ is the polarisation of the optical mode and we have used the definition eq. (3.19). It should be noted that the definition of g used here is half the usual value to avoid an abundance of factors of two.

6.1.1 The interaction picture

Many systems of interest consist of two interacting subsystems. In cases where the subsystems have simple, well-understood dynamics it can be useful to recast the problem in the interaction picture [b55]. In this approach, a hybrid between the Schrödinger picture and the Heisenberg picture is used to model the system. In the Schrödinger picture the state of the quantum system evolves under the influence of the Hamiltonian according to eq. (3.14) while the observables are held constant. This can be recast by incorporating the time evolution into the observables and leaving the state vectors fixed, giving the Heisenberg picture.

In the interaction picture, the evolution of each system independently of the other is treated in the Heisenberg picture while the energy that is associated with the interaction between the systems is considered in the Schrödinger picture. The effect is somewhat analogous to considering a classical mechanics problem in a moving reference frame to simplify the equations of motion.

To move to the interaction picture, the Hamiltonian is decomposed into two parts

$$\hat{H} = \hat{H}_0 + \hat{H}_{int} \tag{6.7}$$

where $\hat{H}_0$ includes the energy of each subsystem in the abscence of any interaction, and $\hat{H}_{int}$ is the interaction energy. We incorporate the evolution of $\hat{H}_0$ into the wavefunction

$$|\tilde{\psi}\rangle = \hat{U}^\dagger|\psi\rangle. \tag{6.8a}$$

$$\hat{U} = e^{-i\hat{H}_0 t/\hbar} \tag{6.8b}$$

and $|\tilde{\psi}\rangle$ now evolves according to $i\hbar|\tilde{\dot{\psi}}\rangle = \hat{\hat{H}}|\tilde{\psi}\rangle$. The interaction picture Hamiltonian $\hat{\hat{H}}$ that governs the evolution of $|\tilde{\psi}\rangle$ is given by

$$\hat{\hat{H}} = \hat{U}^\dagger \hat{H}_{int} \hat{U} \tag{6.9a}$$

$$= -i\hbar \hat{U}^\dagger \hat{\dot{U}} + \hat{U}^\dagger \left(\hat{H}_0 + \hat{H}_{int}\right) \hat{U} \tag{6.9b}$$

Similarly, we can move to any rotating frame by incorporating the rotation into the Hamiltonian. To move to a frame rotating at a frequency ω, we make the transformations

$$|\tilde{\psi}\rangle = \hat{U}^\dagger|\psi\rangle, \qquad \hat{\hat{H}} = -i\hbar \hat{U}^\dagger \hat{\dot{U}} + \hat{U}^\dagger \hat{H} \hat{U}, \qquad \hat{U} = e^{-i\omega t}. \tag{6.10}$$

6.1.2 The rotating wave approximation

We can transform the Hamiltonian into the interaction picture so that the evolution at optical frequencies can be isolated from evolution at the time-scales of the experiment. Decomposing the system Hamiltonian into two parts, $\hat{H}_0 = \hat{H}_{atom} + \hat{H}_{field}$ and $\hat{H}_{int}$, we apply a unitary transformation

$$\hat{U} = e^{-i(\omega_0 \hat{\sigma}_{ee} + \omega \hat{a}^\dagger \hat{a})t}. \tag{6.11}$$

to $\hat{H}_{int}$ to find the interaction picture Hamiltonian

$$\hat{\hat{H}} = \hat{U}^\dagger \hat{H}_{int} \hat{U} \tag{6.12}$$

. Using spectral decomposition of the operator, along with the fact that the atomic and optical subspaces commute, the unitary transformation can be written

$$\hat{U} = \left(\hat{\sigma}_{gg} + e^{-i\omega_0 t}\hat{\sigma}_{ee}\right) \otimes \left(e^{-i\omega \hat{a}^\dagger \hat{a} t}\right). \tag{6.13}$$

Applying $\hat{U}_a = (\hat{\sigma}_{gg} + e^{-i\omega_0 t}\hat{\sigma}_{ee})$ to the atomic subspace of $\hat{H}_{int}$ we find

$$\hat{U}_a^\dagger \left(\hat{\sigma}_+ + \hat{\sigma}_-\right) \hat{U}_a = e^{-i\omega_0 t}\hat{\sigma}_+ + e^{i\omega_0 t}\hat{\sigma}_-, \tag{6.14}$$

and applying $\hat{U}_o = \exp\left(-i\omega\hat{a}^\dagger\hat{a}t\right)$ to the optical subspace of $\hat{H}_{int}$, we get

$$\hat{U}_o^\dagger\left(\hat{a}^\dagger + \hat{a}\right)\hat{U}_o = e^{-i\omega t}\hat{a}^\dagger + e^{i\omega t}\hat{a}. \tag{6.15}$$

For the optical part of the transformation, we used spectral decomposition of the operators, along with the commutators $[\hat{a}^\dagger\hat{a}, \hat{a}^\dagger] = \hat{a}^\dagger$ and $[\hat{a}^\dagger\hat{a}, \hat{a}] = -\hat{a}$ to find the relations

$$e^{-i\omega\hat{a}^\dagger\hat{a}t}\hat{a}e^{i\omega\hat{a}^\dagger\hat{a}t} = e^{i\omega t}\hat{a} \tag{6.16a}$$

$$e^{-i\omega\hat{a}^\dagger\hat{a}t}\hat{a}^\dagger e^{i\omega\hat{a}^\dagger\hat{a}t} = e^{-i\omega t}\hat{a}^\dagger. \tag{6.16b}$$

The the two parts of the interaction picture Hamiltonian can be put together to arrive at

$$\hat{\tilde{H}} = \hbar g\left(e^{-i\omega_0 t}\hat{\sigma}_+ + e^{i\omega_0 t}\hat{\sigma}_-\right) \otimes \left(e^{-i\omega t}\hat{a}^\dagger + e^{i\omega t}\hat{a}\right) \tag{6.17a}$$

$$= \hbar g\left(e^{-i(\omega_0+\omega)t}\hat{\sigma}_+ \otimes \hat{a}^\dagger + e^{-i\Delta t}\hat{\sigma}_- \otimes \hat{a}^\dagger + e^{i\Delta t}\hat{\sigma}_+ \otimes \hat{a} + e^{i(\omega_0+\omega)t}\hat{\sigma}_- \otimes \hat{a}\right) \tag{6.17b}$$

where we have defined $\Delta \equiv \omega - \omega_0$.

The interaction picture Hamiltonian consists of four time-dependent terms, two rotating at frequencies of $\pm\Delta$ and $\pm(\omega_0 + \omega)$. The rotating wave approximation is to neglect the terms that rotate at $\pm(\omega_0 + \omega)$, twice the optical frequency, because they will rapidly average to zero on the time-scales of interest. After making this approximation, we transform back to the Schrödinger picture to obtain the standard Jaynes-Cummings Hamiltonian

$$\hat{\tilde{H}} = \hbar\omega_0\hat{\sigma}_{ee} + \hbar\omega\hat{a}^\dagger\hat{a} + \hbar g\left(\hat{\sigma}_-\hat{a}^\dagger + \hat{\sigma}_+\hat{a}\right). \tag{6.18}$$

The neglected terms correspond to the operators $(\hat{\sigma}_+ \otimes \hat{a}^\dagger)$, the atomic and optical raising operators, and $(\hat{\sigma}_- \otimes \hat{a})$,the atomic and optical lowering operators. These terms therefore correspond to the situations of exciting the atom while gaining a photon and de-exciting the atom while losing a photon. From the perspective of energy conservation we could intuitively expect that contributions from terms would be negligible and, for weak interactions near resonance, they are.

6.1.3 The semi-classical approach

The Jaynes-Cummings model provides a fully quantum treatment of light-atom interaction. For systems that achieve strong light-atom coupling, the quantisation of the optical field influences the dynamics of the atomic evolution. For such systems, the dynamics are found by solving the Schrödinger equation for a state evolving under the action of the Jaynes-Cummings Hamiltonian. The systems of interest in this thesis, however, do not exhibit strong light-atom coupling. For weak coupling, it is sufficient to treat the interaction semi-classically by using a quantum model of the atom, but treat the optical field as classical.

The justification lies partly in the linearity of the equations with respect to the weak optical field and partly in an analysis of the quantum noise that is added to the quantum optical state. It has been shown [a19, a57, a58] that the noise added to the quantum optical state by the decay terms corresponds to the addition of only vacuum noise. The consequence is that decay due to decay processes in the atoms affect the optical state in the same way as passive loss on a beam-splitter. The intuitive reason is that a spontaneous transition from $|g\rangle$ to $|e\rangle$ is energetically unfavourable because the optical frequency is well above the thermal excitation for room-temperature atoms and a transition from $|g\rangle$ to $|s\rangle$ is dipole forbidden [a19]. The decay processes are therefore strictly dissipative.

The Hamiltonian for the semi-classical approach is

$$\hat{H} = \hat{H}_{atom} + \hat{H}_{int}, \tag{6.19a}$$

$$\hat{H}_{atom} = \hbar\omega_0\hat{\sigma}_{ee}; \qquad \hat{H}_{int} = -\hat{\boldsymbol{d}}\cdot\boldsymbol{E}. \tag{6.19b}$$

The interaction part of the Hamiltonian is identical to the fully quantum treatment, minus the hat on $\boldsymbol{E}$. If we assume a form for the electric field of

$$\mathbf{E} = \boldsymbol{\epsilon}E_0\cos(\omega t) = \frac{1}{2}\boldsymbol{\epsilon}E_0\left(e^{i\omega t} + e^{-i\omega t}\right), \tag{6.20}$$

where $\boldsymbol{\epsilon}$ is the polarisation unit vector and E_0 is the real amplitude of the electric field, we can define a quantity

$$\Omega \equiv \frac{1}{2\hbar}(\boldsymbol{d}_{eg}\cdot\boldsymbol{\epsilon})E_0 \tag{6.21}$$

known as the the Rabi frequency. Like the definition of g, the value of Ω used here is half the usual value. With this, the interaction Hamiltonian becomes

$$\hat{H}_{int} = -\hbar\hat{\sigma}_+(\Omega e^{-i\omega t} - \Omega e^{i\omega t}) + h.c. \tag{6.22}$$

We can follow a similar procedure to the quantum case and perform a rotating wave approximation in the interaction picture or, by comparison to eq. (6.5b), identify $\Omega e^{-i\omega t}$ ($\Omega^* e^{i\omega t}$) as being analogous to $g\hat{a}$ ($g\hat{a}^\dagger$) in the quantum case. Following this comparison, we can write the semi-classical Hamiltonian

$$\hat{\hat{H}} = \hbar\omega_0\hat{\sigma}_{ee} - \hbar(\Omega e^{-i\omega t}\hat{\sigma}_+ + \Omega^* e^{i\omega t}\hat{\sigma}_-) \tag{6.23}$$

by analogy to eq. (6.18). The time dependence of the Hamiltonian can be removed by considering the problem in a frame that rotates with the optical field. We apply the unitary transformation, $\hat{U} = e^{i\omega t|e\rangle\langle e|}$ and find that the Hamiltonian in the rotating frame is

$$\hat{\hat{H}} = \hbar\Delta\hat{\sigma}_{ee} - \hbar(\Omega\hat{\sigma}_+ - \Omega^*\hat{\sigma}_-). \tag{6.24}$$

This provides the evolution for an atom driven by an optical field in a compact form with no rapid time-dependence. We can extend this to more complex atoms without any change to the procedure but usually begin by writing the Hamiltonian in the same form as eq. (6.23) directly.

Equation (6.24) provides a model for the evolution of an atom driven by a classical optical field. This semi-classical treatment of the atomic evolution is the approach used for the rest of this thesis, however, when discussing weak fields a notational sleight-of-hand is often performed: the Rabi frequency Ω used to denote a classical field is swapped for the quantum vector potential

$$\Omega \to g\hat{a}. \tag{6.25}$$

For systems with weak light-atom coupling, the semi-classical treatment applies for quantum fields and the result is unaffected, however, it is worth emphasising that the atomic evolution is derived without taking into account the effect of a quantised optical field. The notation is used as a reminder that the optical field could be in a non-classical state.

6.2 Damping and the pure-state approximation

The Schrödinger equation describes the evolution of a pure state, however, the atomic ensemble is subject to sources of decoherence that can be considered to be interaction with a thermal background. This interaction causes information about the system to be lost and the state will become a statistical mixture between the pure state evolution described by the Schrödinger equation and a thermal state.

6.2.1 Mixed states and the density matrix

If a quantum system is in a statistical mixture of several quantum states the density matrix formalism can used to to describe it. In this case, there is a technical randomness associated with the state of the system, as opposed to the intrinsic randomness associated with the probability for a given measurement outcome in quantum mechanics. Here, the randomness is a result of a lack of knowledge of the experimenter with regards to the state of the system. The formalism used to describe this situation is to define a density operator

$$\hat{\rho} = \sum_i p_i |\psi_i\rangle\langle\psi_i| \tag{6.26}$$

where p_i is the probability that the system is in the state $|\psi_i\rangle$. The density matrix, written in terms of an orthonormal basis $\{|u_i\rangle\}$, has elements

$$\rho_{mn} = \langle u_m|\hat{\rho}|u_n\rangle. \tag{6.27}$$

The expectation value for an observable $\hat{O}$ of a mixed state is given by

$$\langle\hat{O}\rangle_\rho = \mathrm{Tr}(\hat{\rho}\hat{O}) \tag{6.28}$$

The density matrix of a pure state $\hat{\rho} = |\psi\rangle\langle\psi|$ evolves according to

$$i\hbar\partial_t\hat{\rho} = [\hat{H}, \hat{\rho}]. \tag{6.29}$$

For a mixed state, however, the evolution is described by a master equation that includes the effect of coupling to a thermal reservoir. We assume a master equation in Lindblad form [b55]

$$\begin{aligned}\partial_t\hat{\rho} &= \frac{1}{i\hbar}[\hat{H}, \hat{\rho}] - \hat{\mathcal{L}}_{decay}(\hat{\rho}), & (6.30)\\ &= \frac{1}{i\hbar}[\hat{H}, \hat{\rho}] - \sum_{j=n}\left(2\hat{J}_n\hat{\rho}\hat{J}_n^\dagger - \hat{\rho}\hat{J}_n^\dagger\hat{J}_n + \hat{J}_n^\dagger\hat{J}_n\hat{\rho}\right), & (6.31)\end{aligned}$$

where $\hat{J}_n$ are operators that describe a spontaneous jump between two states. These jump operators depend on the dynamics of the decay processes and the form that they take in the master equation ensures that the trace of $\hat{\rho}$ is preserved. The only decay mechanism that is modelled in this thesis is population relaxation. For the two-level atom, that means decay from $|e\rangle$ to $|g\rangle$. For this processes, the jump operator is $\hat{J} = \sqrt{\gamma/2}|g\rangle\langle e|$ at a rate γ. This gives $\hat{J}^\dagger\hat{J} = \gamma/2|e\rangle\langle e|$ and

$$\hat{\mathcal{L}}_{decay}(\hat{\rho}) = \frac{\Gamma}{2}\begin{pmatrix} 2\rho_{ee} & \rho_{eg} \\ \rho_{ge} & -2\rho_{ee}\end{pmatrix}. \tag{6.32}$$

Combining eqs. (6.23), (6.30) and (6.32), the evolution of the two-level atom driven by an electric field is

$$\partial_t \rho = \begin{pmatrix} -\Gamma\rho_{ee} + i(\rho_{ge}\Omega + \rho_{eg}\Omega^*) & -(i\Delta + \frac{\Gamma}{2})\rho_{eg} + i\Omega(\rho_{gg} - \rho_{ee}) \\ (i\Delta - \frac{1}{2}\Gamma)\rho_{ge} + i\Omega^*(\rho_{gg} - \rho_{ee}) & \Gamma\rho_{ee} - i(\rho_{ge}\Omega + \rho_{eg}\Omega^*) \end{pmatrix} \tag{6.33}$$

To simplify solving the evolution of the system, it is common to assume that the ensemble lies almost entirely in the ground state which, for weak driving fields, is well justified. In this approximation, known as the pure-state approximation we assume that

$$\hat{\sigma}_{gg} \approx 1, \qquad \hat{\sigma}_{ee} \approx 0. \tag{6.34}$$

In this approximation we the atomic state entirely by the coherence $\hat{\sigma}_{ge}$:

$$\partial_t \hat{\sigma}_{eg} = -(i\Delta + \Gamma/2)\hat{\sigma}_{eg} + i\Omega. \tag{6.35}$$

6.2.2 Collective excitation

With a quantum model for the dynamics of a damped atom driven by an optical field, the next step is to determine the effect that an ensemble of atoms will have on an optical field. To do this, we introduce collective operators that describe the overall state of the ensemble. If we assume that the ensemble contains enough atoms that we can approximate it as a continuous distribution of thin slices, each containing $N_z \gg 1$ atoms, we can define collective operators for the atoms contained in each slice:

$$\hat{\sigma}_{\mu\mu}(z,t) = \frac{1}{N_z}\sum_{j=1}^{N_z} \hat{\sigma}^j_{\mu\mu}(t) \tag{6.36}$$

$$\hat{\sigma}_{eg}(z,t) = \frac{1}{N_z}\sum_{j=1}^{N_z} \hat{\sigma}^j_{es}(t) e^{-i\omega(t - z_j/c)} \tag{6.37}$$

$$\tag{6.38}$$

The bulk polarisation resulting from the ensemble will be

$$\hat{\mathcal{P}} = \eta(\boldsymbol{d}_{eg}\cdot\boldsymbol{\epsilon})\hat{\sigma}_{ge} \tag{6.39}$$

where η is the atomic density of the ensemble.

Similarly to eq. (3.19) we define a slowly-varying operator

$$\hat{\mathcal{E}}(z) = \sqrt{\frac{2c}{\Gamma L}} e^{i\omega_0(t - z/c)} \int_\omega \hat{a}_\omega e^{i\omega z/c}, \tag{6.40}$$

where L is the length of the ensemble, by including the explicit time-dependence of the optical frequency in the definition. The dimensionless normalisation ensures simple equations of motion.

6.2.3 The equations of motion

The propagation of an optical pulse through a material is given by eq. (3.30). We can combine this with the polarisation of the ensemble, eq. (6.39) and the relation between the electric field and the vector potential, eq. (3.19). We arrive at the propagation equation for an optical pulse travelling through an atomic ensemble

$$(\partial_t + c\partial_z)\,\hat{\mathcal{E}} = igN\sqrt{\frac{2c}{\Gamma L}}\hat{\sigma}_+ \tag{6.41}$$

assuming a uniform distribution of N atoms, or

$$(\partial_t + c\partial_z)\,\hat{\mathcal{E}} = ig\eta(z)L\sqrt{\frac{2c}{\Gamma L}}\hat{\sigma}_+ \tag{6.42}$$

where $\eta(z)$ is the linear atomic density, otherwise.

The equations of motion can be further simplified by considering a moving reference frame eq. (3.31) and rescaling some of the quantities [a57]. We introduce the definition $\hat{P} = \sqrt{N}\hat{\sigma}_{eg}$ and the dimensionless coordinate

$$\xi(z) = \int_0^z dz'\eta(z')/N \tag{6.43}$$

which ranges from 0 to 1 and reduces to $\partial\xi = \partial z/L$ when $\eta(z)$ is a constant value over the length of the ensemble. In the transformed coordinates, eqs. (6.35) and (6.42) for the rescaled operators are

$$\partial_\tau \hat{P}(\xi,\tau) = -\left(i\Delta + \frac{\Gamma}{2}\right)\hat{P}(\xi,\tau) + i\sqrt{d}\frac{\Gamma}{2}\hat{\mathcal{E}}(\xi,\tau) \tag{6.44a}$$

$$\partial_\xi \hat{\mathcal{E}}(\xi,\tau) = i\sqrt{d}\hat{P}(\xi,\tau). \tag{6.44b}$$

The parameter $d = 2g^2NL/(\Gamma c)$ is the optical depth of the atomic ensemble on resonance. The intensity of a weak monochromatic optical field propagating through the ensemble will be attenuated by a factor of e^{-2d}.

It is relatively uncommon to scale the spatial coordinate to be dimensionless while the temporal coordinate is left in real units. Dividing eq. (6.44a) by Γ renders the equations of motion dimensionless and simplifies the expression. The reason that I prefer to leave the temporal coordinates in real units is that it is practical to perform calculations directly in terms of frequencies that can be measured in the lab. This simplifies translation between simulation results and experiments. The exact location of an optical within an ensemble however, tends not to be of interest in the lab because only the input and output fields are accessible. The equations of motion as written here, eqs. (6.44), are entirely in terms of quantities that are directly accessible experimentally: detuning, linewidth, and optical depth.

6.2.4 Solutions in the Fourier domain

We can solve for the transmission through the ensemble by taking a temporal Fourier transform of the operators defined by $\hat{a}(\tau) = \int_\omega \hat{a}(\omega)e^{-i\omega\tau}$. In the Fourier domain, $\partial_\tau \to -i\omega$, and, taking $\Delta = 0$, we combine eqs. (6.44a) and (6.44b) to arrive at the propagation equation

$$\partial_\xi \hat{\mathcal{E}}(\omega,\xi) = \frac{d\Gamma}{2(i\omega - \Gamma/2)}\hat{\mathcal{E}}(\omega,\xi). \tag{6.45}$$

From eqs. (3.33) and (6.45) we can identify

$$\chi(\omega) = \frac{i}{\tilde{k}} \frac{d\Gamma}{(i\omega - \Gamma/2)} \tag{6.46}$$

as the linear susceptibility of the ensemble for the weak field. Here, $\tilde{k}$ is the wave-number in the dimensionless coordinates ξ and, for a uniform distribution of atoms, $\tilde{k} = kL$. The solution to eq. (6.45) is then

$$\hat{\mathcal{E}}(\omega, \xi) = e^{ik\chi(\omega)\xi} \hat{\mathcal{E}}(\omega, 0) \tag{6.47}$$

and we can define the transfer function that describes the impulse response for propagation through the atomic ensemble as

$$\hat{\mathcal{E}}(t, \xi = 1) = \int_{\omega} T(\omega, \xi = 1) \hat{\mathcal{E}}(\omega, \xi = 0) e^{i\omega t} d\omega, \tag{6.48a}$$

$$T(\omega, \xi = 1) = e^{i\chi(\omega)}. \tag{6.48b}$$

The real and imaginary parts of of the electric susceptibility $\chi(\omega)$ are responsible for the dispersion and absorption of $\hat{\mathcal{E}}(\omega)$, respectively. Note that when $\omega = 0$ the *amplitude* of the incident field is attenuated by e^{-d}.

6.3 The driven three-level atom

The topics in this thesis on quantum memory use ensembles of three-level atoms that are driven by a weak field that drives a transition between $|g\rangle$ and $|e\rangle$ as well as by a strong field that drives a transition between $|e\rangle$ and a meta-stable state $|s\rangle$. The level structure, including detunings and decay rates is shown in fig. 6.1. Using the same methods and approximations as for the two-level atom, we can find the equations of motion for the three-level atom starting from the Hamiltonian

$$\hat{H} = \hat{H}_0 + \hat{H}_{int} \tag{6.49}$$

$$\hat{H}_0 = \hbar\omega_{se}\hat{\sigma}_{ss} + \hbar\omega_{ge}\hat{\sigma}_{ee} \tag{6.50}$$

$$\hat{H}_{int} = -\hbar \left(\Omega\hat{\sigma}_{se} + g\sqrt{\frac{L}{2\pi c}} \int d\omega \hat{a}_{\omega} \hat{\sigma}_{ge} \right) + H.C. \tag{6.51}$$

where Ω is the Rabi frequency of the bright field with frequency ω_{se}, $\hat{a}_{\omega}$ is a weak quantum field and $H.C.$ is shorthand for the Hermitian conjugate. For this system, we assume that the bright coupling field drives all of the population to lie in $|g\rangle$ such that

$$\hat{\sigma}_{gg} \approx 1, \quad \hat{\sigma}_{ee} \approx 0, \quad \hat{\sigma}_{ss} \approx 0, \quad \hat{\sigma}_{es} \approx 0. \tag{6.52}$$

In the rotating wave, pure-state and slowly-varying envelope approximations the equations of motion for the system are then

$$\partial_t \hat{\sigma}_{eg} = -(i\Delta + \Gamma)\hat{\sigma}_{eg} + ig\hat{\mathcal{E}} + i\Omega\hat{\sigma}_{gs} \tag{6.53}$$

$$\partial_t \hat{\sigma}_{es} = -(i\delta + \gamma)\hat{\sigma}_{es} + i\Omega^* \hat{\sigma}_{ge} \tag{6.54}$$

$$(\partial_t + c\partial_z)\hat{\mathcal{E}} = ig\eta(z)L\hat{\sigma}_{eg} \tag{6.55}$$

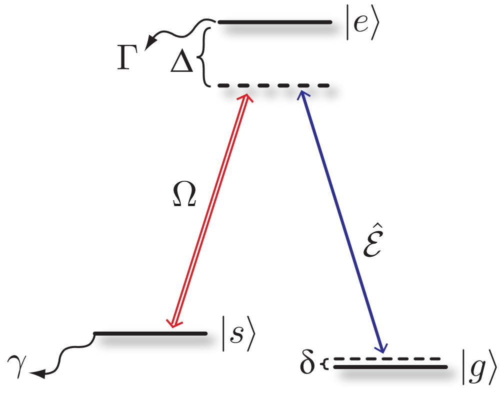

Figure 6.1: The energy level structure of the Λ-type three-level atom.

where γ is the population decay rate from $\hat{\sigma}_{ss}$ to $\hat{\sigma}_{gg}$ and $\eta(z)$ is the linear atomic density. The factor of two has been omitted from the decay rates $\Gamma/2 \to \Gamma$ and $\gamma/2 \to \gamma$ for simplicity. This implies that the rates refer to the dephasing rates of the excited and metastable states, respectively, as opposed to the population decay. With this, the definition of the optical depth becomes $d = g^2NL/(\Gamma c)$. The detuning δ is the frequency difference between the coupling and signal fields relative to the frequency difference between $|g\rangle$ and $|s\rangle$. We can simplify the equations using the definitions $\xi(z) = \int_0^z dz'\eta(z')/N$, $\tau = t - z/c$, $\hat{P} = \sqrt{N}\hat{\sigma}_{eg}$, $\hat{S} = \sqrt{N}\hat{\sigma}_{es}$, and $\hat{\mathcal{E}} \to \sqrt{c/(L\Gamma)}\hat{\mathcal{E}}$ to arrive at

$$\partial_\tau \hat{P} = -\left(\Gamma + i\Delta\right)\hat{P} + i\sqrt{d}\Gamma\hat{\mathcal{E}} + i\Omega\hat{S}, \tag{6.56a}$$

$$\partial_\tau \hat{S} = -\left(\gamma + i\delta\right)\hat{S} + i\Omega^*\hat{P}, \tag{6.56b}$$

$$\partial_\xi \hat{\mathcal{E}} = i\sqrt{d}\hat{P}. \tag{6.56c}$$

Expressing the equations of motion in terms of the optical depth of the ensemble, and thus the dimensionless coordinate ξ, means that the system can be modelled entirely in terms of parameters that are directly measurable in an experiment. This removes the dependence on g and N, both of which are difficult to determine. In contrast, quantities that have units of frequency, such as dephasing rates or detunings, are comparatively straightforward to measure in the lab and the use of a normalised temporal unit would only obscure the relation between the equations and the experiment.

As for the case of the two-level atom, we solve the equations in the Fourier domain, $\partial_\tau \to -i\omega$, and obtain

$$\hat{P} = i\frac{\sqrt{d}\Gamma\hat{\mathcal{E}} + \Omega\hat{S}}{\Gamma + i(\Delta - \omega)}, \tag{6.57a}$$

$$\hat{S} = i\frac{\Omega^*\hat{P}}{\gamma + i(\delta - \omega)}. \tag{6.57b}$$

Combining eq. (6.56c) with eqs. (6.57a) and (6.57b) and taking $\delta = 0$ we arrive at the propagation equation

$$\partial_\xi \hat{\mathcal{E}}(\omega, \xi) = \chi(\omega)\hat{\mathcal{E}}(\omega, \xi) \tag{6.58}$$

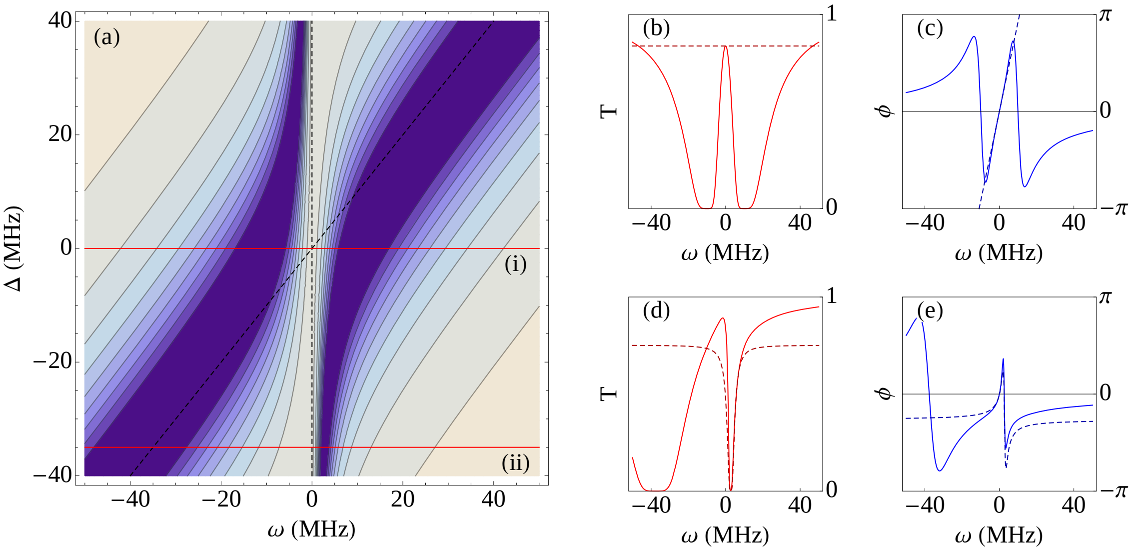

Figure 6.2: (a) The transmittance eq. (6.60a) through an ensemble of driven three level atoms for varying one- and two-photon detunings. The black dashed lines indicate $\omega = 0$ and $\omega = \Delta$. The plots (b-e) show the transmittance (b,d) and phase shift (c,e) through the ensemble at one-photon detunings of $\Delta = 0$ MHz (b,c) and $\Delta = 35$ MHz (d,e). The dashed lines show the approximated functions using adiabatic elimination of the excited state for the resonant and off-resonant cases.

where the dispersion is given by

$$\chi = d\Gamma\left(i(\omega-\Delta) - \Gamma + \frac{|\Omega|^2}{i\omega-\gamma}\right)^{-1}. \tag{6.59}$$

From eqs. (6.48) we can see that the transmittance T and phase shift ϕ through the ensemble will be

$$T(\omega) = e^{-2\mathbf{Im}[\chi(\omega)]} \tag{6.60a}$$

$$\phi(\omega) = \mathbf{Re}[\chi(\omega)]. \tag{6.60b}$$

Figure 6.2 (a) shows the transmittance as a function of Δ and δ for an ensemble with the parameters $d = 5$, $\Gamma = 6$, $\Omega = 10$, and $\gamma = 0.3$. The figure illustrates a number of important features of a driven three-level atom. When $\Delta \to 0$, the effect of electromagnetically induced transparency is apparent in the transparency window that is opened in the centre of the atomic absorption feature. The transmittance and phase shift for this case are shown on the right in the plots (b, c).

The off-resonant case is illustrated by the plots of transmittance and phase shift marked (d, e) which correspond to a detuning of $\Delta = 35$ MHz. These show the Raman absorption line that results from the two-photon process. The ac-Stark shift of the one- and two-photon resonances is visible in the shift of the resonances relative to the black dashed lines that indicate $\omega = 0$ and $\omega = \Delta$.

6.4 Adiabatic approximations

The equations of motion, eqs. (6.56), can be simplified by making the assumption that the slowly varying envelope for the excited state coherence, $\hat{P}$, varies adiabatically. The term adiabatic means that the state is in an instantaneous eigenstate of the Hamiltonian. In this context, the condition requires that $\partial_t \hat{P} \ll (\Gamma + i\Delta)$. This is satisfied when the evolution of the excited is dominated by the intrinsic damping rate Γ or the detuning Δ.

There are two distinct limiting cases for adiabatic elimination: the resonant case and the far off-resonant case. The resonant case applies in the limit of

$$\Delta \ll d\Gamma \tag{6.61}$$

while the off-resonant case requires

$$\Delta \gg d\Gamma. \tag{6.62}$$

6.4.1 The resonant limit

We examine the resonant case by setting $\Delta \to 0$ in eq. (6.59) and take a power series expansion to the first order in ω to arrive at

$$\chi(\omega) \approx \frac{d\Gamma}{\gamma\Gamma + |\Omega|^2}\left(i\gamma + \frac{\gamma^2 + |\Omega|^2}{\gamma\Gamma + |\Omega|^2}\omega + \mathcal{O}(\omega^2)\right). \tag{6.63}$$

This provides the susceptibility in the limit of slowly varying signal fields that are in two-photon resonance with the pump field, which is in turn in one-photon resonance with the atoms. The adiabatic approximation applies in the limit of

$$\gamma + i\omega \ll \frac{|\Omega|^2}{\Gamma}, \tag{6.64}$$

which will be discussed in greater detail in chapter 8 along with the implications for pulse propagation in EIT. The adiabatic approximation is illustrated in fig. 6.2 (b,d) by the dashed lines. The zeroth order approximation provides a constant value for the absorption while the first order approximation provides the dispersion.

6.4.2 The Raman limit

In the far-detuned regime, the evolution of the excited state is dominated by the detuning. In this case we can adiabatically eliminate the excited state using the assumption that the characteristic time-scale of the interaction is small relative to the detuning Δ. We apply this assumption when taking a Fourier transform of eq. (6.56a), and let $\Delta - \omega \approx \Delta$ so that

$$\hat{P} = i\left(\sqrt{d}\Gamma\hat{\mathcal{E}} + \Omega\hat{S}\right)/\left(\Gamma + i\Delta\right) \tag{6.65}$$

Inserting this into Equations (6.56b) and (6.56c) we arrive at

$$\partial_\tau \hat{S} = -(\gamma' + i\delta')\hat{S} + i\sqrt{d}\Gamma\frac{\Omega^*}{\tilde{\Delta}}\hat{\mathcal{E}} \tag{6.66a}$$

$$\partial_\xi \hat{\mathcal{E}} = i\left(d\frac{\Gamma}{\tilde{\Delta}}\hat{\mathcal{E}} + \sqrt{d}\frac{\Omega}{\tilde{\Delta}}\hat{S}\right) \tag{6.66b}$$

where

$$\delta' = \delta - \Delta\frac{|\Omega|^2}{\Gamma^2+\Delta^2}, \quad \gamma' = \gamma + \Gamma\frac{|\Omega|^2}{\Gamma^2+\Delta^2}, \quad \tilde{\Delta}^{-1} \equiv \frac{\Delta + i\Gamma}{\Gamma^2+\Delta^2}. \tag{6.67}$$

The modified values δ' is the detuning adjusted to include the ac-Stark shift induced by the bright coupling field and, similarly, γ' is the ground-state dephasing rate adjusted to account for the coupling field power-broadening. The real and imaginary parts of the term $-d(\Gamma/\tilde{\Delta})\hat{\mathcal{E}}$ are the absorption and dispersion, respectively, due to the excited state. For negligible absorption, we require that $d\mathbf{Re}[\Gamma/\tilde{\Delta}] \ll 1$. This gives the limit for the Raman regime, $\Delta/\Gamma \gg d$. While the dispersion due to the detuned excited state is small in the Raman regime, it can still be included in the equations of motion by incorporating it into the envelopes for $\hat{\mathcal{E}}$ and Ω by setting $\hat{\mathcal{E}} \to \hat{\mathcal{E}}e^{id\xi/(1+\Delta^2/\Gamma^2)}$ and $\Omega \to \Omega e^{id\xi/(1+\Delta^2/\Gamma^2)}$. With this, and treating the off-resonant absorption as negligible, the equations of motion become

$$\partial_\tau \hat{S} = -(\gamma' + i\delta')\hat{S} + i\sqrt{d}\Gamma\frac{\Omega^*}{\tilde{\Delta}}\hat{\mathcal{E}} \tag{6.68a}$$

$$\partial_\xi \hat{\mathcal{E}} = i\sqrt{d}\frac{\Omega}{\tilde{\Delta}}\hat{S}. \tag{6.68b}$$

For large detunings, it is reasonable to take $\tilde{\Delta} \simeq \Delta$. This approximation is is illustrated in fig. 6.2 (c,e) by the dashed lines. The approximated susceptibility is that of an equivalent two-level atomic ensemble as described by eqs. (6.68) added to the off-resonant susceptibility of the ensemble in the absence of a pump field.

The final simplified set of equations resembles the equations of motion for a two-level atom with an effective light-atom coupling rate $g\frac{\Omega}{\Delta}$. We can complete the analogy to a two level atom by defining $\hat{\hat{S}} \equiv (\Omega/\Delta)\hat{S}$. With this rescaling, we can write

$$\partial_\tau \hat{\hat{S}} = -(\gamma' + i\delta')\hat{\hat{S}} + i\sqrt{d}\gamma'\hat{\mathcal{E}} \tag{6.69a}$$

$$\partial_\xi \hat{\mathcal{E}} = i\sqrt{d}\hat{\hat{S}}. \tag{6.69b}$$

where we have made the approximations

$$\tilde{\Delta} \simeq \Delta \text{ and } \gamma' \approx \Gamma\left(\frac{|\Omega|}{\Delta}\right)^2. \tag{6.70}$$

These approximations are valid in the limit of very large detuning relative to the natural linewidth and in the limit where the Raman linewidth is dominated by the power broadening of the pump field rather than the decay rate of $\hat{S}$. The equations are identical to those of a two-level atom but, importantly for proposals in this thesis, the amplitude of $\hat{\hat{S}}$ is optically tunable via the amplitude of the pump field.

6.5 Numerical modeling

The equations of motion that govern the propagation of light in atomic ensembles can be solved numerically to investigate how the system is expected to respond to input pulses. For this task, we use the software XMDS [a59, a60], an acronym for *eXtensible multi-dimensional simulator*. This software constructs C++ code to solve systems of ordinary differential equations (ODEs) that have been described in the form of an XML document. The XML description of the problem includes information regarding how the equations should be solved, such as the numerical integration method and the grid dimensions of the simulation.

The software provides a great deal of flexibility when it comes to defining methods for solving ODEs, such as a variety of internal integration algorithms, stochastic methods, and parallelism. For problems regarding the propagation of light in atomic ensembles, however, we use only relatively simple methods. All of the problems discussed in this thesis are simplified to consider propagation in only one spatial dimension and use an even two-dimensional grid. A method-of-lines approach is used to solve the system of equations with a boundary value condition applied at one end of the propagation dimension. The temporal dimension is solved using an adaptive fourth-fifth order Runge-Kutta numerical integration, while the spatial dimension is solved using a non-adaptive fourth-order Runge-Kutta method.

Using the XMDS solver, the memory dynamics can be simulated to provide an intuitive understanding of the behaviour and to predict and optimise efficiencies. Appendix A.2 provides sample code for a simple simulation of a two-level gradient echo memory.

CHAPTER 7

Experimental implementation

The design and construction for both the warm-atom and cold-atom memory experiments was carried out primarily by my colleagues. My aim for describing the experimental set-ups is therefore to provide context for the experimental results contained in this work, as opposed to a complete characterisation of the methods. The warm-atom experiment was initially developed by Gabriel Hétet and Mahdi Hosseini and a thorough description of the set-up is contained in [a49, c50]. The cold-atom memory was built by Ben Sparkes, Julien Bernu, and Nick Robins and the project is detailed in [a5, c52].

7.1 Rubidium

All of the experiments were conducted using rubidium for the atomic ensembles. It is an alkali metal with a relatively low melting point of 39°C. These properties give it two of the important characteristics for experiments in quantum optics: it has a single valence electron, providing a simple energy level structure, and it has a high vapour pressure at relatively low temperatures. The two common isotopes are ^{85}Rb and ^{87}Rb with nuclear spins of 5/2 and 3/2, respectively. The nuclear spins provide a rich hyperfine electronic energy level structure, as shown in figs. 7.1 and 7.2, that can be exploited for experiments in quantum optics. The final attractive property of rubidium is that the optical transitions to the lowest excited states is addressable by inexpensive diode lasers so that even complex optical schemes can be implemented at a relatively low cost.

Some of the properties rubidium that need to be commonly referred to for calculations and simulations are included here in table 7.1 and figs. 7.1 and 7.2. These values are taken from [a61] and [a62] for ^{85}Rb and ^{87}Rb, which contain much more complete summaries of the physical properties of rubidium than is included here.

Property	Value for ^{85}Rb	Value for ^{87}Rb	Unit
$5^2S_{1/2} \rightarrow 5^2P_{1/2}$ (D$_1$) transition	794.98		nm
$5^2S_{1/2} \rightarrow 5^2P_{3/2}$ (D$_2$) transition	780.21	780.24	nm
Line-width (D$_1$)	$2\pi \times 5.75$		MHz
Line-width (D$_2$)	$2\pi \times 6.07$		MHz

Table 7.1: Common physical parameters of rubidium

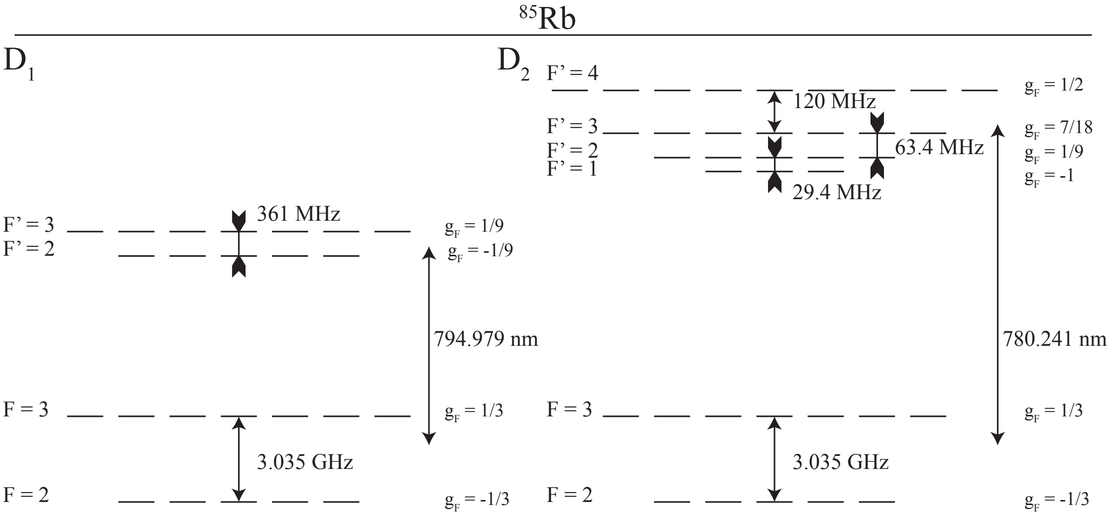

Figure 7.1: Level structure of ^{85}Rb. The most useful frequencies and Landé g-factors are shown.

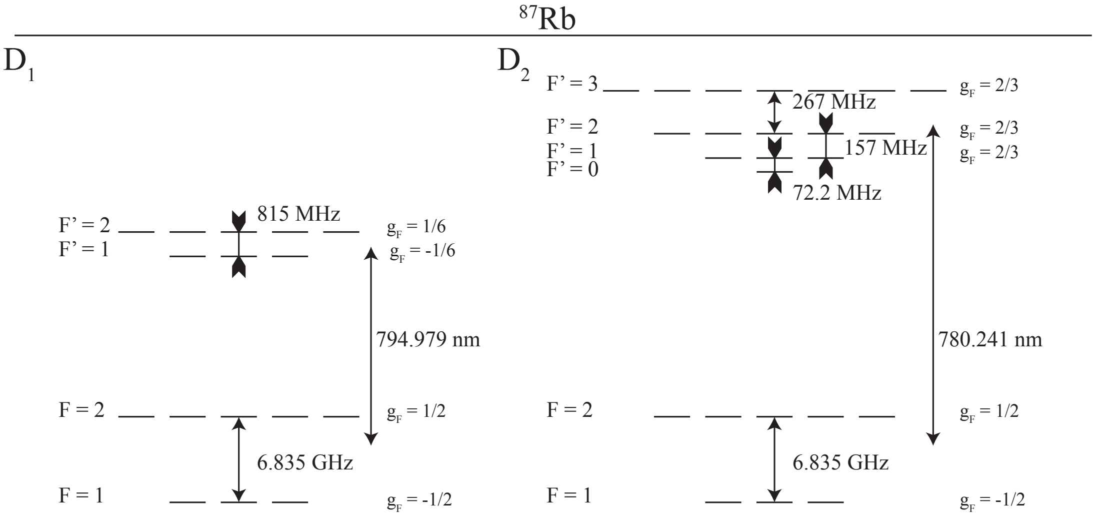

Figure 7.2: Level structure of ^{87}Rb. The most useful frequencies and Landé g-factors are shown.

Many of techniques related to coherent light-atom interaction in rubidium use the Zeeman effect to manipulate the energy level structure. The Zeeman effect induces an energy level shift that is proportional to the magnetic field and to the z-component of the total angular momentum of the atom. Explicitly, in the limit where the effect of the magnetic field is small compared to the hyperfine splitting, the Zeeman splitting is [a62]

$$\Delta E = \mu_B g_F m_F B \tag{7.1}$$

where μ_B is the Bohr mangneton and g_F is the Landé g-factor. The splitting of the Zeeman sub-levels with the application of a magnetic field is one of the main tools used for the gradient echo memory as well as for neutral atom trapping.

7.2 Experimental techniques

7.2.1 The lasers

Experiments on rubidium require lasers that operate at either 780 nm or 795 nm. Ti:Sapphire lasers are capable of addressing either of these transitions and produce a high power, up to 3.5 W, while maintaining a narrow linewidth. These lasers are pumped by a single-mode 1064 nm laser that has been frequency doubled to 532 nm. Both of the experimental set-ups used for experiments covered in this thesis used these to provide the probe and control fields for the memory.

Two brands of Ti:Sapphire were used. The earlier experiments used a Coherent MBR laser, providing about 1.5 W of optical power, while later experiments switched to an M^2 SolsTiS laser, providing about 2 W. Both lasers can be tuned over the rubidium D_1 and D_2 transitions. The Coherent MBR was replaced due to a 87 kHz modulation that was generated within the laser to lock the intra-cavity etalon. The M^2 lasers uses a similar locking technique, but features a substantially suppressed modulation at the output relative to the Coherent system.

Additional lasers used in each experiment, such as the trapping lasers for the MOT or lasers for additional optical pumping, are diode lasers due to the lower cost. These produce up to 100 mW of light with a linewidth of around 200 kHz. They consist of a laser diode and a diffraction grating that retro-reflects the first diffraction order into the diode to form an external cavity.

The quality of the output light depends on the type of diode used. Cheap diode that are not anti-reflection (AR) coated are less stable and more prone to mode-hopping than the AR coated variety. The lasers themselves are provided by MOGlabs, Toptica, or are home built. The MOGlabs diode lasers tend to be relatively stable and are used for the trapping and re-pump lasers in the cold-atom experiment.

The high powers required for the trapping laser in the cold-atom experiment are achieved by amplifying a diode laser with a tapered amplifier. These amplifiers contain a diode slab that increases in width along the beam propagation axis. These diodes are AR coated and, when seeded by tens of milliwatts of laser power, will amplify it to up to a watt of optical power. The tapered amplifiers that we use contain diodes that are produced by eagleyard Photonics and are mounted in temperature stabilised home-made mountings.

7.2.2 Frequency locking

In order to form a magneto-optical trap, the trapping and re-pump lasers must be frequency locked to atomic transitions. This is performed using a dither locking technique similar to Pound-Drever-Hall locking. Instead of phase-modulating the laser, however, the frequency of rubidium atoms in a reference cell is modulated by applying an oscillating magnetic field to the atomic sample. The absorption of the sample will modulate the intensity of the laser, which can then be demodulated using a lock-in amplifier, or even a simple mixer and low-pass filter, to provide an error signal.

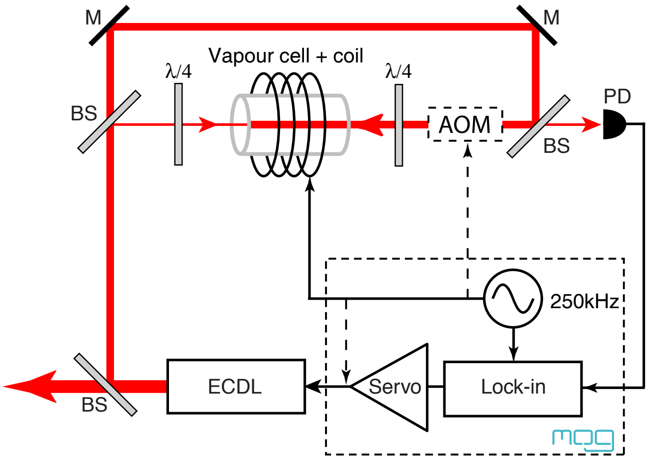

Figure 7.3: The set-up for saturated absorption locking. The image is taken from the MOGlabs laser manual.

To improve the locking accuracy, a Doppler-free saturated absorption spectrum is used. This is done by illuminating the reference cell with a counter-propagating pump field, as shown in fig. 7.3. The signal field and the counter-propagating pump each address atoms with opposite velocities depending on the detuning from the center of the Doppler-broadened linewidth. Only when the laser is tuned in the center of the Doppler linewidth do the signal and pump fields address the same atoms, resulting in a transparency feature for the signal that is as wide as the natural linewidth. It is this feature to which the laser is locked to provide a stable reference.

Further flexibility in the laser frequency can be gained by placing an acousto-optic modulator between the laser and the atomic reference cell. This allows a frequency-shifted beam to be locked to an atomic transition, providing a frequency offset for the laser. Using this technique, a diode laser can be locked anywhere within the hyperfine splitting of the D_2 excited states.

7.2.3 Generating the probe field

The signal fields for the experiments are produced by a fibre-coupled electro-optic phase-modulator followed by a locked cavity. The modulator generates side-bands at ≈ 6.8 GHz relative to the coupling field. The cavity is used to select only the correct side-band, as shown is fig. 7.4. A additional EOM is used to place an additional modulation on the beam for PDH locking.

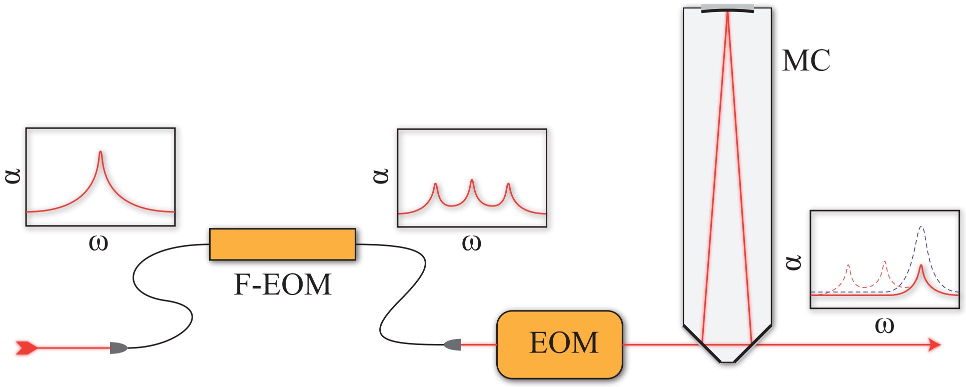

Figure 7.4: A probe field for the memory experiments must be generated with a frequency difference of one Rb ground-state splitting relative to the pump. This is done using a fibre-coupled wave-guide electro-optic-optic modulator. A mode-cleaning cavity is then used to pass only the side-band that is useful for the probe field. A second EOM is used to produce the modulation needed to lock this cavity.

7.3 GEM experiments in hot atoms

Experiments on atomic ensemble based memories are most easily carried out using warm vapours of alkali metals such as rubidium or caesium. The basic set-up for a light-atom interaction experiment in these metals is a simple, sealed glass cell. To improve overall efficiencies and reduce scatter, the glass cell is usually anti-reflection coated and is also magnetically shielded from the environment. The magnetic shielding is critical to ensure that the quantization axis of the atoms can be defined along the optical path. Additional magnetic fields parallel to the optic axis may be introduced according to the memory protocol.

The shielding consists of a hollow cylinder of μ-metal that extends well past the ends of the vapour cell. The shielding substantially reduces the effect of the Earth's magnetic field as well as any local magnetic sources. This reduction is crucial because any magnetic fields that are transverse to the optical axis will broaden the intrinsic linewidth of the atomic Zeeman sub-levels. It is worth noting, however, that the large uniform magnetic fields that are applied along the optical axis to split the Zeeman sub-levels mean that small external magnetic fields that are parallel to the optical axis do not reduce to performance of the memory.

7.3.1 Merging the control and signal fields

The control and signal fields need to be combined in the same polarisation for the highest efficiency operation of the memory. For classical fields, this can be done with a beamsplitter, however the combination is necessarily inefficient as one output port from the beamsplitter is unused. For quantum states light loss in the system must be reduced to a minimum in order to preserve the states. To achieve efficient combination of the beams, a ring-cavity was used. The control field was transmitted while the signal field was reflected into the same mode, as shown in fig. 7.6.

The cavity was locked using a counter-propagating field locked to the TEM_{01} mode to avoid contaminating the signal beam or adding modulation to the control. An AOM was

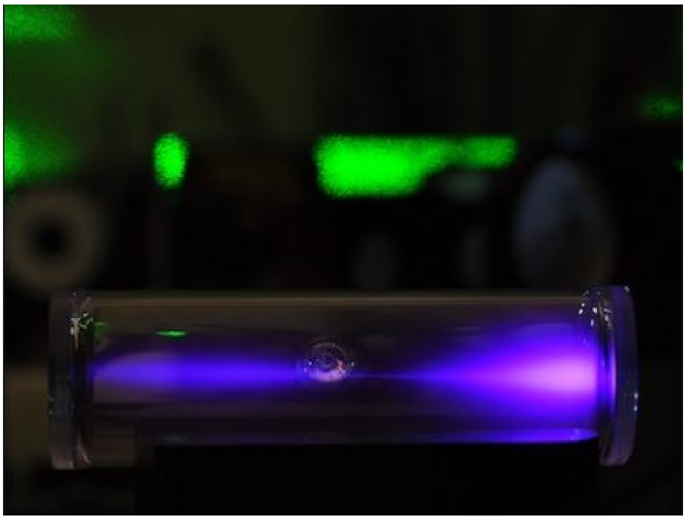

Figure 7.5: Fluorescence from a rubidium vapour cell that is being illuminated with a focused resonant beam. The purple colour is due to dipole-dipole interactions between excited rubidium atoms [a63]

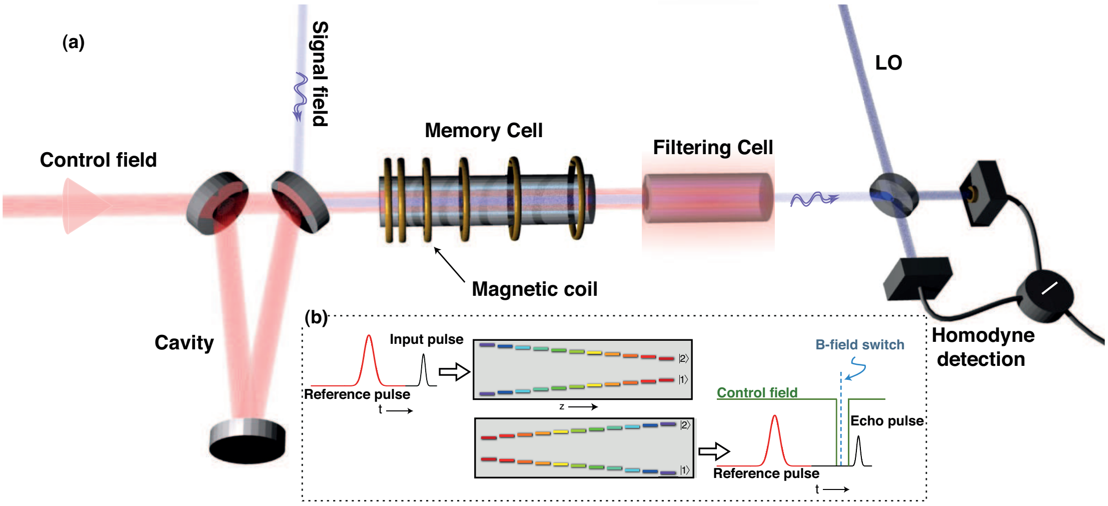

Figure 7.6: A schematic of the optical path for warm-atom memory experiments. An optical cavity is used to combine the control and signal field in the same polarisation before the memory. After the memory, a filtering cell containing ^{85}Rb is used to absorb the control field so that the weak signal field can be detected. The magnetic field coils are shown as loops to illustrate generation of the gradient.

used to frequency shift this locking beam so that the control field would be resonant with the TEM_{00} mode while locked.

7.3.2 The vapour cell

The vapour cells used are 20 cm long and are filled with rubidium that has been isotopicaly enhanced to contain a higher ratio of ^{87}Rb than is present in natural abundance. The optical depth required to obtain a high efficiency memory is obtained by heating the rubidium vapour cell to increase the partial pressure of rubidium vapour. The cell is heated to approximately 80°C using wrappings of resistive heating wire. A solid-state relay switch turns off the heaters synchronously with the experiment so that they do not introduce additional magnetic fields.

The lifetime of the ground-state coherence depends on a number of decay mechanisms. The time-of-flight through the optical mode would be the limiting factor if not for the use of a buffer gas to slow the diffusion rate of rubidium atoms in the cell. Krypton buffer gas at a pressure of 1 Torr is used. The time-of-flight decay is further reduced by using large diameter beams of about 1 cm.

7.3.3 Control field filtering

In order to measure the weak signal beam, elimination of the control field is required. Two approaches to accomplish this were used. The first was to fibre couple the probe field to provide strong spatial filter. Only a minimal amount of control field light is also coupled into the fibre because it has a significantly larger beams waist and a different divergence to the signal field. For the experiments to measure the quantum state, however, this filtering is still inadequate and, in any case, the high control field power risks damaging the fibre.

The solution found was to employ a second vapour cell containing rubidium with an isotopic ratio of natural abundance. The ^{85}Rb that is present in this cell has a transition that is conveniently on resonance with the control field but not the signal field. At temperatures of $\approx 100°$ C, the optical depth of the ^{85}Rb is sufficient to provide an extinction ratio of more than 60 dB. Interestingly, a bright violet fluorescence is observed in the filtering cell when the control field is incident. The cause is a collisional energy transfer that populates higher excited stated that can then emit radiation [a63]. The effect is shown in fig. 7.5.

7.3.4 Gradient coils

The magnetic fields for GEM are produced by three magnetic field coils contained within the μ-metal magnetic shielding. A solenoid produces a uniform field that splits the Zeeman levels and two gradient coils broaden the Raman lines corresponding to magnetically sensitive transitions. Each gradient coil is wound with a variable pitch from one end of the cell to the other, producing a linearly increasing magnetic field along the propagation axis. The coils have pitches that vary in opposite directions. When a current is switched from one coil to the other, the effect is that the sign of the gradient changes while the offset of the magnetic fields from zero remains roughly the same. The current is switched using solid-state relay switches and power resistors are used in series with the coils to decrease the switching time.

7.4 GEM experiments in cold atoms

To improve the storage time that is possible in the memory, experiments were also conducted with an ensemble of cold atoms that are prepared using a magneto-optical trap. The magneto-optical trap employs cooling and trapping methods to slow down and confine neutral atoms to a small volume. The low temperatures reduce the Dopper broadening to well under the natural linewidth of the atoms, allowing long coherence times and no inhomogeneous broadening. Furthermore, the large atomic densities that are possible with trapping provide the high optical depths needed for memory experiments.

The cold atoms are trapped in the centre of a glass chamber that is attached to a larger stainless steel vacuum system. The vacuum system is continuously pumped out with an ion pump to maintain pressures of roughly 10^{-10} Torr. Rubidium is introduced to the system by dispensers, produced by SAES Group, that contain rubidium salts. These release atomic rubidium when a heating current is passed through them. The dispensers have a lifetime that is limited by the quantity of rubidium that they contain. For the SAES sources, this was found to be six months to a year of daily use. When the dispensers were exhausted, we experimented with using dispensers from Alvatec. These are based on rubidium alloys rather than salts and have a larger quantity of rubidium. After a number trials with the Alvatec dispensers, however, we had repeated problems with contamination from the dispensers and decided to continue using models from SAES.

7.4.1 Doppler cooling

The trap uses optical Doppler cooling to slow rubidium atoms within a vacuum chamber. The cooling is done using red-detuned, circularly polarised beams that illuminate the atoms from all directions. The red-detuned photons of the trapping field have a lower energy than the atomic resonance. An cycle consisting of the absorption of a trapping photon followed by spontaneous emission results in a loss of energy from the atom into the radiation field. This energy comes from the kinetic energy of the atom, resulting in a cooling effect.

In terms of momentum transfer, atoms are more likely to absorb a photon if they are moving in the opposite direction to the photon due to the Doppler shift. While the emission is in a random direction, the momentum change due to the absorption of a photon always opposes the velocity of the atom. The result in an effective viscosity, known as optical molasses.

7.4.2 Magneto-optical trapping

The trapping is accomplished by applying a magnetic field gradient to the trapping region. This gradient is tuned to have a zero in the center of the trap. Atoms that are away from the point of zero magnetic field experience a Zeeman shift that makes them more likely to absorb a photon that provides momentum back towards the center of the trap. Correctly chosen circular polarisations ensure that the sign of the Zeeman shift is such that atoms are more likely to absorb a photon that directs them back to the centre of the trap.

The combination of Doppler cooling and magneto-optical trapping provide both the cooling mechanism and the trapping mechanism with one laser field. To be effective, the atoms must cycle through absorption and emission of the trapping light. There is a chance, however, that each spontaneous emission will decay into the hyperfine ground state that is not addressed by the trapping laser. A re-pump laser is therefore used to pump atoms back into hyperfine state that is acted on by the trapping fields.

The trapping laser is 30 MHz red-detuned from the D_2 $F = 3 \rightarrow F' = 4$ transition and the re-pump is resonant with the D_2 $F = 2 \rightarrow F' = 3$ transition. The re-pump and trapping beams are collimated and combined before being delivered to the trapping region. A dark spot in the centre of the re-pump can help to increase the number of trapped atoms by allowing the atoms that are already in the centre to remain out of the trapping cycle, reducing heating by radiation pressure. The dark spot was effective when used with a three-dimensional trapping field, but seemed to be counter-productive for a two-dimensional trap.

7.4.3 The trap geometry

The trapping region is elongated by using race-track shaped magnetic field coils. The elongated trap serves to match the trapped volume to the laser beam as well as possible. The optical depth scales proportionally to g^2NL, where $g \propto 1/\sqrt{V}$ is the light-atom coupling rate, L is the length of the ensemble, and N is the number of atoms in the interaction volume V. The optical depth is therefore proportional to the total number of atoms divided by the cross-sectional area of the ensemble so, it is advantageous to use a long, narrow trap.

The trap can either use weak axial magnetic field coils to confine the atoms in the direction of propagation or the axial fields can be left off to form a two-dimensional trap. Both configurations have produced optical depths of around 1000 on the $F = 2 \rightarrow F' = 3$ transition in ^{85}Rb. The cross-section of the trap is further reduced using a compression step prior to performing experiments. We used a 480 ms load time to collect a large number of atoms. The MOT is compressed in two dimensions by increasing the transverse magnetic field gradients along with the detunings of the trapping and re-pump fields over 20 ms [a5]. The signal field is focused along the long axis of the atom cloud, with a beam waist of 200 μm.

7.4.4 Control field filtering

The near lack of Doppler broadening that is present in the cold atom system allows a more straightforward approach to filtering out the control field than the hot atom system because a small angle can be introduced between the control field and the signal field. The control beam was collimated to a diameter of 7 mm and overlaps the signal field where the MOT is located. We employ spatial filtering in two stages. The first stage is a razor blade which blocks the majority of the control, the second is a pinhole through which the signal is focused. The spatial filtering provides an extinction ratio of $\approx$ 45 dB for the control while maintaining $\approx$ 80% efficiency for the signal.

7.4.5 Gradient magnetic fields

The magnetic fields used for GEM experiments are produced by Helmholtz coils, oriented to provide a magnetic field along the direction of propagation. Each of the coils is driven by an independent bi-polar current source. The current common to the two coils provides an offset magnetic field to split the Zeeman sub-levels while the current difference applies a magnetic field gradient across the ensemble. The current sources are powered by PA107 op-amps from Apex Power and can provide up to 5 amps and 100 volts with a rise-time of less than 1 μs.

7.4.6 Preparation of the ensemble

After the compression phase, the magnetic field and re-pump laser are turned off. The trapping laser is left on to continue Doppler cooling the atoms and to prepare the ensemble in the $F = 1$ hyperfine ground state. Because memory experiments require precise control of the magnetic field that is applied to the ensemble, the magnetic fields used for trapping must be switched off prior to performing the experiment. Eddy currents induced in the optical bench are allowed to die away for 1 ms after the trapping fields are removed before the memory experiment begins. The coupling field is usually turned on 400 μs before the writing stage to further pump the atoms to the $F = 1$ hyperfine ground state.

7.4.7 Memory experiments with cold atoms

As with the hot atom experiment, the coupling field and two signal fields are derived from a Ti:Sapphire laser. For GEM experiments, the laser is set 200 MHz blue-detuned from the ^{87}Rb D$_1$ $F = 2 \rightarrow F' = 2$ transition. All three fields pass through separate acousto-optic modulators to allow for fine frequency adjustment as well as gating and pulsing.

The signal fields are combined with the same linear polarisation (horizontal), passed through an optical fibre and focused to a waist that matches that of the atomic cloud (200μm). The coupling field is polarised orthogonally to the signal fields and is collimated to a larger diameter of 7 mm so that it illuminates the entire ensemble uniformly and propagates with a small angle relative to the signal beams. The experimental set-up is schematically shown in fig. 7.7.

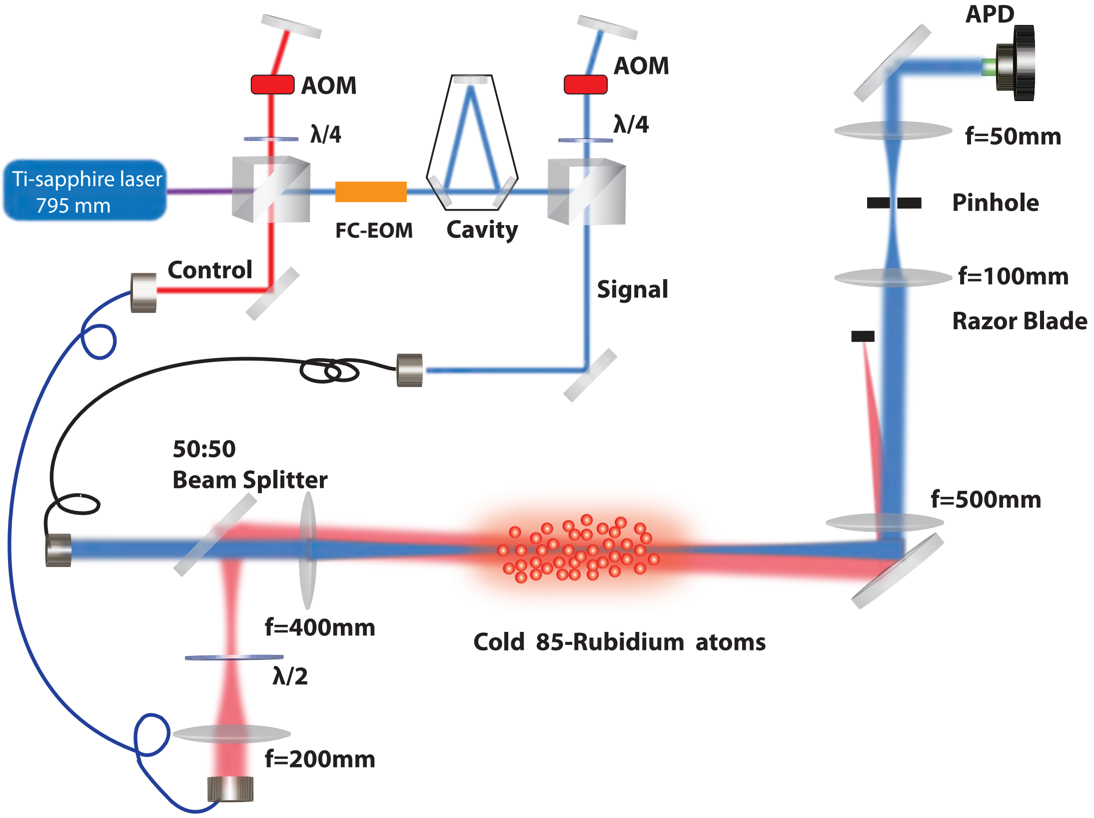

Figure 7.7: A schematic of the cold-atom experiment. The experiment is similar to the warm-atom experiment with the exception of the combination and separation of the control and signal fields. The small Doppler broadening of the cold-atom experiment allows a small angle to be introduced between the fields so that they can be spatially separated.

CHAPTER 8

EIT and 4WM

Electromagnetically induced transparency has been a major avenue of investigation for research into quantum memories for over a decade. Its relative simplicity makes it an excellent platform for proof-of-principle experiments and the first demonstrations of storage and recall of quantum states of light were done using EIT. The performance of EIT is directly related to the optical depth of the ensemble. A larger optical depth provides a larger dispersion and therefore a lower group velocity for optical pulses propagating through the ensemble. For efficient memories, group delays that are larger than the pulse width are required to store the entire pulse and, for temporally multi-mode memories, the delay must be sufficient to store multiple input pulses. The delay, however, increases with the square-root of the optical depth, requiring extremely high optical depths, upwards of ≈ 200, for good memory performance.

At large optical depths, however, four-wave mixing (4WM) becomes significant and gain during slow-light propagation begins to contribute to the output [a64]. In this case, noise is necessarily added to the output [a65] and the performance of EIT as a quantum memory is degraded. Experimental work in warm atomic vapours has demonstrated significant 4WM [a66, a67], however, cold-atom memories are of interest due to the high storage efficiencies that have been reported [a35]. The un-broadened excited state available in cold atom systems also yields a lower 4WM strength than warm vapour EIT for identical optical depth, which, in principle, allows a high storage-efficiency with little added noise.

Experiments that use warm vapours as the atomic ensemble can achieve the large optical depths required for EIT memories, however, the high temperatures result in diffusion of the atoms and collisions that degrade the memory performance. Experiments in cold atoms reduce these effects and reduce the Doppler broadening to less than the natural linewidth. They are therefore excellent platforms for implementing EIT-based memories but obtaining a sufficiently large optical depth has been a challenge. The high-optical depth cold atom cloud developed for experiments on gradient echo memories, however, has a sufficient optical depth for excellent EIT memory performance and we took the opportunity to investigate EIT in high OD cold atomic ensembles and to quantify the role of four-wave-mixing in such systems.

Previous work [a35] has shown that high-performance memory is possible in large OD cold-atomic clouds but the optical depths used in those experiments was at the limit where four-wave-mixing is beginning to contribute to the output. Using still higher optical depths and comparing EIT in both common isotopes of rubidium, we demonstrate that 4WM has a non-negligible contribution to the intensity of the signal field after EIT delay for sufficiently large optical depths. Good agreement between our results and theory indicate that the theoretical model used in [a65] is a realistic description of cold-atom EIT. We also show that with sufficiently high optical depth, delays of more than two pulse-widths are possible, thus allowing temporally multi-mode delay. Previous experiments have demonstrated large fractional delays using the strong dispersion of dense optical vapours [a68], however, to our knowledge this is the first example of EIT with sufficient delay to store multi-mode pulses.

8.1 The conditions for slow light

Before discussing the role of four-wave-mixing and the approach that we take to modelling it, I will introduce the dynamics of EIT slow-light and storage. The equations of motion that govern the propagation of a weak signal pulse under the conditions of EIT are a special case of eqs. (6.56a) to (6.56c), where the detuning is zero and the pump field Ω is larger than the ground-state decay rate. The susceptibility obtained from these equations in the Fourier domain is given by eq. (6.59) with the transmittance T and phase shift ϕ given by eqs. (6.60), as shown in fig. 8.1.

For EIT, we consider a signal field near two-photon resonance with the pump and take a power series expansion to the first order in ω to arrive at an approximation, eq. (6.63). The zero-th and first order terms give approximations of imaginary and real parts of the susceptibility, respectively. In this limit, we find

$$T \approx \exp\left(-\frac{2d\Gamma}{\Gamma + |\Omega|^2/\gamma}\right) \tag{8.1a}$$

$$\phi(\omega) \approx -d\Gamma\omega\frac{\gamma^2 + |\Omega|^2}{(\gamma\Gamma + |\Omega|^2)^2}. \tag{8.1b}$$

These approximations are shown on Figure 6.2 (b,c) as dashed lines.

The reduction in absorption from $\exp -2d$ is the result of the atomic ensemble being prepared in a dark-state, in which the excited state has no contribution. This dark-state can be intuitively understood by looking at eq. (6.56a) when where $\Delta = 0$ and in the limit where $\partial_t \hat{P} \to 0$. Under these conditions, the excited state is given by

$$\hat{P} = i\sqrt{d}\hat{\mathcal{E}} + i\frac{\Omega}{\Gamma}\hat{S}. \tag{8.2}$$

If

$$\hat{S} = -\sqrt{d}\Gamma\hat{\mathcal{E}}/\Omega \tag{8.3}$$

then $\hat{P}$ remains unexcited and there is no absorption of the signal field. This is known as the dark-state. If the system begins in the dark-state, for example if there is initially no signal field, and evolves sufficiently slowly, it will remain in the dark-state even as the signal field increases. This is adiabatic evolution, which is when the system remains in an eigenstate as the Hamiltonian changes.

The condition for adiabatic evolution is given by eq. (6.64). If this condition is satisfied along with $\gamma\Gamma \ll |\Omega|^2$, the optical depth of the ensemble is reduced by a factor of $\gamma\Gamma/|\Omega|^2$.

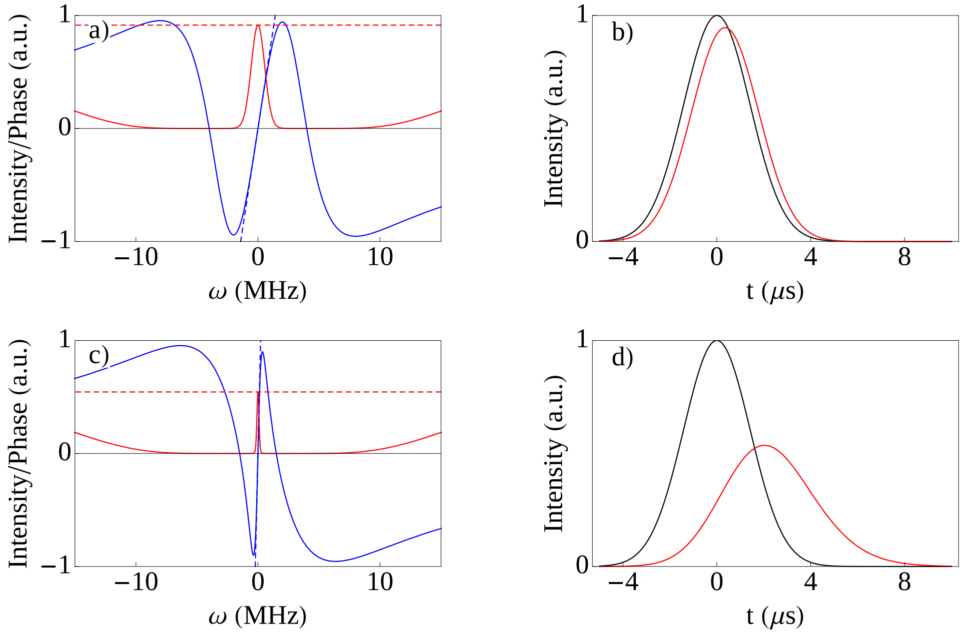

Figure 8.1: (a,c) The transmittance (red) and dispersion (blue) from the real and imaginary parts of the susceptibility, eqs. (8.1), respectively. From the transfer function for the system, the response to an input pulse can be calculated. (b,c) show the output from the EIT system (red) when excited by an input pulse (black) that is centred about zero detuning. The pump field Rabi frequencies are 4 MHz for (a,b) and 1.5 MHz for (c,d). The other parameters used are $d = 6$, $\Gamma = 6$, $\gamma = 0.02$

The conditions for negligible absorption are then

$$\frac{2d\Gamma}{|\Omega|^2} \ll \tau_c, \text{ and } \frac{2d\Gamma}{|\Omega|^2} \ll \gamma^{-1}, \tag{8.4}$$

where τ_c is the temporal width of the pulse that is sent through the ensemble. The slow-light aspect of EIT can be seen from the group velocity on resonance

$$v_g \equiv \frac{d\omega}{dk} = c\left[n(\omega) + \omega_0\left(\frac{dn}{d\omega}\right)\right]^{-1} \tag{8.5}$$

$$\approx \left[\frac{1}{c} + \frac{d\Gamma}{L}\frac{\gamma^2+\Omega^2}{(\gamma\Gamma+|\Omega|^2)^2}\right]^{-1} \tag{8.6}$$

where $n(\omega) \approx 1 + \mathbf{Re}[\chi(\omega)/2]$ and ω_0 is the optical frequency. For $|\Omega|^2 \gg \{\gamma\Gamma, \gamma^2\}$, this simplifies to

$$v_g \approx \left[\frac{1}{c} + \frac{d\Gamma}{|\Omega|^2 L}\right]^{-1}. \tag{8.7}$$

Figure 8.1 (b,d) illustrates this slowing of an optical pulse. The plots show the output pulse (red) along with the input pulse (black) as a reference. The pulses are delayed, with a larger delay being observed for the smaller pump power, and thus narrower transparency feature and steeper dispersion.

For efficient storage, we require that the entire pulse be contained within the ensemble. This can usually only happen for group velocities that are much lower than c because the pulse durations, on the order of a few microseconds, correspond to pulses that are kilometres long in free space while the vapour cells or cold atom clouds are on the order of centimetres. The condition for the pulse to be contained within the ensemble is that the time-bandwidth product must be greater than one. For group velocities much less than c we can approximate the time delay as

$$\Delta\tau = \frac{L}{v_g} \approx d\frac{\Gamma}{|\Omega|^2}. \tag{8.8}$$

To estimate the bandwidth of the transparency window ω_c we can consider the second order term in the expansion eq. (6.63) and find that, in the limit of $\gamma \to 0$, the bandwidth is given by

$$\omega_c \propto \frac{|\Omega|^2}{\sqrt{d}\Gamma} \tag{8.9}$$

resulting in a time-bandwidth product $\Delta\tau \cdot \omega_c \propto \sqrt{d}$. For a high-efficiency memory, we require an optical depth $\sqrt{d} \gg 1$.

8.2 Storage and recall

In section 8.1 the susceptibility for a weak probe field propagating through an ensemble of three-level atoms was found under the condition of the atoms being driven on-resonance by a control field that couples the excited state to an auxiliary meta-stable state. The absorption of this driven system shows a transparency window when the frequency difference between the control field and the probe field is equal to the frequency difference between the ground and meta-stable states. This is known as electromagnetically induced transparency and the corresponding dispersion gives rise to a slow light effect. In [a34] it was shown that photons propagating through the slow-light medium and the spin-polarisation of the medium itself can be combined into a quasi-particle: the dark-state polariton. The polariton is a linear superposition of a photon and a spin excitation in the ensemble. The term dark-state refers to the fact that the excited state does not contribute to the polariton and the new quasi-particle is therefore not subject to loss due to spontaneous emission, despite the fact that the pohotonic part is resonant with the excited state transition. A more complete description can be found in [b55] but I will provide a brief derivation of the polariton here.

We can start from the equations of motion for the driven three-level atom, eqs. (6.56a) to (6.56c), and make the assumption that the excited state $\hat{P}$ is not driven by assuming that the conditions eq. (8.4) are satisfied. This places restrictions on the bandwidth of the optical pulse that can be converted into a slow-light polariton efficiently. With these assumptions, eq. (6.56a) becomes

$$\sqrt{d}\hat{\mathcal{E}} + (\Omega/\Gamma)\hat{S} = 0. \tag{8.10}$$

We can define the modes

$$\hat{\Phi} = \sin\theta\hat{\mathcal{E}} + \cos\theta\hat{S} \tag{8.11}$$

$$\hat{\Psi} = \cos\theta\hat{\mathcal{E}} - \sin\theta\hat{S} \tag{8.12}$$

where

$$\sin\theta = \frac{\sqrt{d}}{\beta}, \quad \beta = \sqrt{d + \Omega^2/\Gamma^2}, \quad \cos\theta = \frac{\Omega}{\Gamma\beta}, \tag{8.13}$$

and from eq. (8.10) we find that

$$\hat{\mathcal{E}} = \cos\theta\hat{\Psi}, \quad \hat{S} = -\sin\theta\hat{\Psi}. \tag{8.14}$$

For simplicity, we assume that the control field envelope Ω is real. The mode $\hat{\Phi}$ corresponds to the bright-state and is not excited provided that the assumptions associated with adiabatic elimination of the excited are satisfied. We can infer that the orthogonal mode $\hat{\Psi}$ carries all of the information initially carried by the optical mode and corresponds to the dark-state of the light-atom system. It is therefore known as the dark-state polariton.

We derive the equations of motion for $\hat{\Psi}$ under the assumption that the decoherence γ of the meta-stable state and the two-photon detuning δ are both zero. The equations of motion take the simplest form under the real coordinates z and t as opposed to the transformed coordinates. Assuming a uniform distribution of atoms in the ensemble, eq. (6.56c) becomes

$$\frac{1}{\Gamma}\left(\partial_t + c\partial_z\right)\hat{\mathcal{E}} = i\sqrt{d}\hat{P} \tag{8.15}$$

which can be combined with eq. (6.56b) to give

$$\left(\partial_t + c\partial_z\right)\hat{\mathcal{E}} = \Gamma\frac{\sqrt{d}}{\Omega^*}\partial_t\hat{S}. \tag{8.16}$$

Substituting in eq. (8.14) and noting that the temporal dependence of θ gives $\dot{\theta} = -\partial_\tau\Omega\sin\theta/(\beta\Gamma)$ results in the equations of motion for the system

$$\left(\partial_t + c\cos^2\theta\partial_z\right)\hat{\Psi} = 0, \tag{8.17a}$$

$$\hat{\Phi} = 0. \tag{8.17b}$$

We can identify eq. (8.17a) as an advection equation with a velocity $v_g = c\cos^2\theta$, describing a pulse-preserving propagation at a group velocity that can be controlled by Ω. The dynamics of the polariton are the basis of the idea to EIT as an optical quantum memory. Provided that the adiabatic condition is satisfied and the excited state makes a negligible contribution to the state of the system it is possible to reversibly map information from the optical field to a stationary optical coherence and back again.

In this picture, the incoming light forms a polariton as it enters the atomic ensemble. The dynamics of the polariton can be controlled through the intensity of the applied control field such that it can be entirely converted to a stationary spin coherence and then returned to a propagating optical field. The nature of the memory is inherently single-mode with the optical field being adiabatically converted into the polariton as it enters the ensemble and restored to an optical field as it exits. The mode orthogonal to the dark-state polariton, the bright-state polariton, is negligibly excited, but any contribution from it leads to loss.

8.3 Experiments on EIT in a high optical depth ensemble

The efficiency of an EIT-based quantum memory is directly related to the optical depth, specifically the delay-bandwidth, and thus the efficiency, is proportional to $\sqrt{d}$. A very large optical depth is therefore needed to obtain high-efficiencies in an EIT memory. We meet this challenge by using a magneto-optical trap to confine a large number of cold atoms to a small interaction region, elongated along the optical path of signal field. The

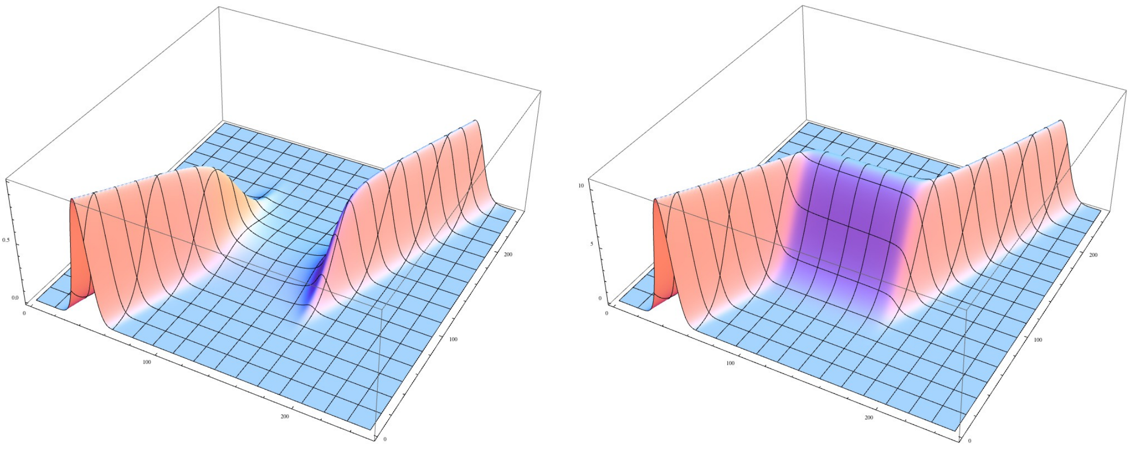

Figure 8.2: Simulation of the EIT polariton.

trapping set-up, described in section 7.4 and presented in much greater detail in [a5], allows the trapping of either ^{85}Rb or ^{87}Rb. For the EIT experiments, ^{85}Rb was primarily used due to a higher optical depth.

To explore the high-optical depth MOT as a platform for an EIT-based quantum memory, we examined the propagation of signal pulses through the atomic ensemble under the conditions of slow light. We illuminate the cold atom cloud with a control field that is substantially larger in diameter than the atomic cloud to ensure a uniform pump field Rabi frequency over the ensemble. The signal pulse is tuned to be exactly on two-photon resonance with the pump and is shaped into a square temporal profile with a 6 μs width. The signal pulses that emerge after propagation through the ensemble were recorded for a variety of control field powers. A reference trace of the signal pulse shape was then recorded by blocking the trapping beams so that the MOT was dispersed. A sample of the results is shown in fig. 8.5 (b). One of the features of the output pulses for short delay times is that they exhibit a small gain rather than loss through the ensemble, as is apparent in the top-most plot of fig. 8.5 (b). The effect is the result of four-wave-mixing, which produces a parametric amplification of the signal and generates a conjugate idler field. In order to analyse the results of EIT slow-light with very high optical depths, we develop a model for the process that includes four-wave-mixing.

8.4 Four-wave mixing

The process of four-wave-mixing is a non-linear optical effect in which two pump photons scatter off an atomic transition to become a signal and an idler photon each with different energies than the pump photons. In the case of four-wave-mixing in EIT, the two pump photons are degenerate, both coming from the control field, and the process is stimulated by the presence of signal photons. This stimulated process leads to coherent gain of the signal, as well as production of an idler. The effect is detrimental to the use of EIT as a quantum memory due to the inherent addition of noise associated with an amplification process. The four-wave-mixing is the result of the strong pump driving the ground state to

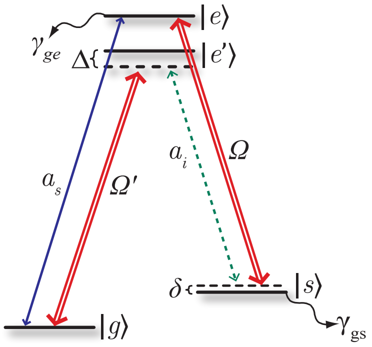

Figure 8.3: The level structure used for the three-level model. The Λ-system for EIT is formed by the control field, Ω and the signal field, a_s. A second Λ-system is formed by Ω′ and the idler field, a_i, which completes a 4WM process.

excited state transition off-resonantly, completing a closed cycle back to the meta-stable state through the generated idler, as illustrated in fig. 8.3.

We assume a four-level atom driven by four optical fields as our model of the system, shown in fig. 8.3. A strong coupling field, with Rabi frequency Ω, couples a meta-stable state, $|s\rangle$, to the excited state $|e\rangle$ and weak signal field, a_s, couples the ground state, $|g\rangle$ to $|e\rangle$, completing a two-photon transition between $|g\rangle$ and $|s\rangle$. the coupling field also couples $|g\rangle$, to $|e'\rangle$ off-resonantly, completing a Raman transition with an idler field a_i which is generated by the four-wave mixing process. Adiabatically eliminating $|e'\rangle$, the equations of motion for this system are [a65]

$$i\partial_t\hat{P} = -i\gamma_e\hat{P} - g_s\sqrt{N}\hat{\mathcal{E}}_s - \Omega\hat{S} \tag{8.18a}$$

$$i\partial_t\hat{S} = -i\gamma_s\hat{S} - g_i\sqrt{N}(\Omega'/\Delta)\hat{\mathcal{E}}_i^\dagger - \Omega^*\hat{P} \tag{8.18b}$$

$$(\partial_t + c\partial_z)\hat{\mathcal{E}}_s = ig_s\sqrt{N}\hat{P} \tag{8.18c}$$

$$(\partial_t + c\partial_z)\hat{\mathcal{E}}_i^\dagger = -ig_i\sqrt{N}(\Omega'/\Delta)\hat{S}, \tag{8.18d}$$

where $\hat{P}$ and $\hat{S}$ are collective spin-polarisation operators corresponding to the $|g\rangle \to |e\rangle$ and $|g\rangle \to |s\rangle$ transitions, respectively, in an ensemble of N atoms. For a three-level atom, Δ is the frequency difference between $|g\rangle$ and $|s\rangle$, although this may be be different for an atom with an additional excited state $|e'\rangle$, as shown in fig. 8.3. The coupling rate for $\hat{\mathcal{E}}_{s(i)}$ is $g_{s(i)}$ and, similarly, $\Omega(')$ is the Rabi frequency for the control field driving $|s\rangle \to |e\rangle$ ($|g\rangle \to |e\rangle$). The rate g_s differs from g_i only due to the different dipole transition strengths associated with each transition, as does Ω from Ω′. γ_e and γ_s are the decay rates from $\hat{P}$ and $\hat{S}$. We solve eq. (8.18a) in the Fourier domain to obtain the expression for the transfer function of the signal field in the absence of an idler:

$$T_s(\omega) = \exp\left[-\frac{d\gamma_e}{4V(\omega)}\left(i\omega - i\omega|\epsilon|^2 + |\epsilon|^2\gamma_e - \gamma_s\right)\right] \\ \times\left(\frac{(\gamma_e|\epsilon|^2 - i\omega - i|\epsilon|^2\omega + \gamma_s)}{U(\omega)}\sinh\left[\frac{D\gamma_e U(\omega)}{4V(\omega)}\right] + \cosh\left[\frac{D\gamma_e U(\omega)}{4V(\omega)}\right]\right), \tag{8.19a}$$

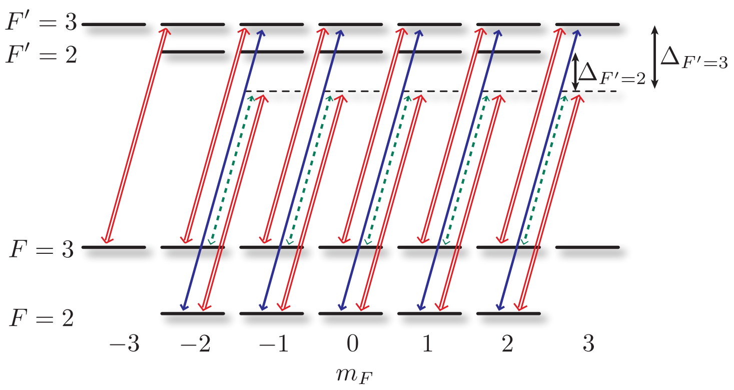

Figure 8.4: The ^{85}Rb level structure used for the experiment showing five degenerate EIT-4WM systems. Each system differs by the Clebsch-Gordan coefficients associated with each transition. The field labelling is shown in fig. 8.3

where

$$U(\omega) = \sqrt{(i\omega + (i\omega - \gamma_{ge})|\epsilon|^2 - \gamma_{gs})^2 + 4|\epsilon\Omega|^2} \tag{8.20a}$$

$$V(\omega) = (i\gamma_{gs} + \omega)(\omega + i\gamma_{ge}) - |\Omega|^2, \tag{8.20b}$$

with the definitions

$$d = 2\frac{g^2NL}{\gamma_e}; \qquad \epsilon = \eta\frac{\Omega}{\Delta}; \qquad \eta = \frac{g_i\Omega'}{g_s\Omega} = \frac{d_{|s\rangle\to|e'\rangle} \cdot d_{|g\rangle\to|e'\rangle}}{d_{|g\rangle\to|e\rangle} \cdot d_{|s\rangle\to|e\rangle}}. \tag{8.21}$$

The parameter η can be expressed in terms of the dipole matrix elements, $d_{|j\rangle\to|k\rangle}$, for the associated transition and the definition of optical depth, d, corresponds to an intensity attenuation of e^{-d} when $\omega \to 0$ and $\Omega \to 0$.

The level structure of the ^{85}Rb atoms that is relevant for the EIT experiment is shown in fig. 8.4. We can identify five degenerate four-level structures coupling from each of the Zeeman sub-levels on the $F = 2$ manifold. These differ only by the Clebsh-Gordan coefficients of the transitions. For each of the degenerate systems, both of the excited state manifolds ($|F' = 2\rangle$ and $|F' = 3\rangle$) are taken into account by summing the strengths of the off-resonant interaction with each excited state, weighted by the relative detuning:

$$\eta_{m_F} = \frac{d_{3,2} \cdot d_{2,2}}{d_{2,3} \cdot d_{3,3}} + \left(\frac{\Delta_{F'=2}}{\Delta_{F'=3}}\right)\frac{d_{3,3} \cdot d_{2,3}}{d_{2,3} \cdot d_{3,3}} = \frac{d_{3,2} \cdot d_{2,2}}{d_{2,3} \cdot d_{3,3}} + \frac{\Delta_{F'=2}}{\Delta_{F'=3}}, \tag{8.22}$$

where $d_{i,j} \equiv d_{|F=i,m_f\rangle\to|F'=j,m_f+1\rangle}$ and $\Delta_{F'=j}$ is the detuning of the idler from the $F' = j$ excited state.

The degenerate EIT systems are reduced to a simple four-level model by the introduction of an effective 4WM strength η_{eff}. We assume that the population is uniformly distributed across the ground-state manifold, in which case η_{eff} can be approximated as the mean of the values of η_{m_F} for each EIT system. We find that for ^{85}Rb, $\eta_{\text{eff}} = 1.62$ and for ^{87}Rb, $\eta_{\text{eff}} = 1.33$.

8.5 Results

The results are analysed by applying the transfer function derived from the four-level model, eqs. (8.19), to the recorded reference trace. A sample of traces recorded at an optical depth of (550 ± 20) is shown in Figure 8.5 (b) along with the theoretical output pulses. The values for optical depths used in the model were obtained using independent off-resonance absorption measurements. The only free parameters are the decay rate, $\gamma_{gs} = (12.8 \pm 0.5)$ kHz, and the ratio between $|\Omega|^2$ and the measured control field intensity which are fitted globally and consistent across all of the data. Other parameters can be found in Ref. [a61] with $\gamma_{ge} = \pi \cdot 5.75\text{MHz}$ and $\Delta = 2\pi \cdot 3.035\text{GHz}$.

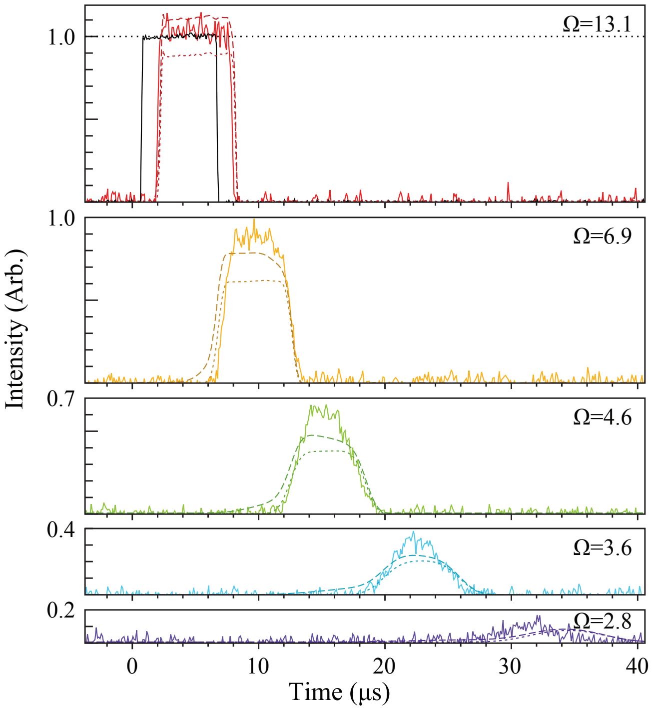

Figure 8.5: Delayed pulses (solid, coloured traces) shown relative to the input (black). Theoretical outputs, obtained by applying eqs. (8.19) to the reference are shown as dashed lines. For comparison, theoretical outputs with no 4WM are shown as dotted lines.

Figure 8.6 shows the integrated output pulse intensities and delay times for a variety of control field powers at different optical depths. The total EIT efficiency and delay times are in good agreement with the model across the entire parameter space that we explored. The theoretical EIT efficiency corresponding to zero 4WM strength are included as a dashed line and the shaded region indicates the effect of 4WM on the efficiency of EIT. At lower optical depths, the theoretical predictions for both cases are indistinguishable with the resolution of the plot. The observed increase in efficiency is in good agreement with the simple four-level model, and contrasts with previous cold atom experiments that have speculated that four-wave mixing reduces the efficiency [a69]. At the highest optical depth achieved in our experiment, where $\text{OD} = 550 \pm 20$, corresponding to a regime where 4WM contributes significantly to the output signal field and additional noise photons would have a detrimental effect on the fidelity of a quantum memory.

In contrast, an experiment conducted with ^{87}Rb in the paper of Chen *et al.* [a35] reported high storage efficiencies with negligible 4WM at an optical depth of 156. This

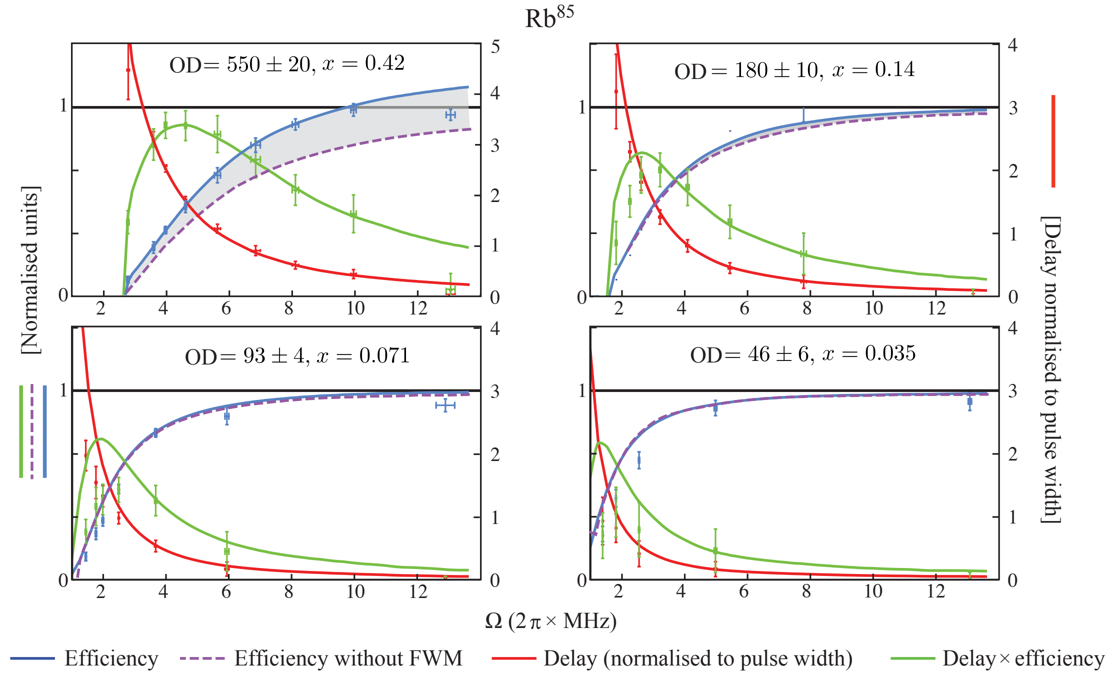

Figure 8.6: The efficiencies (blue) and delay times (red) for ^{85}Rb. The delay times are normalised to the 6 μs input pulse width. The product of delay and efficiency (green) is relative to the reference input. The solid lines show the predicted values based on eqs. (8.19). The dashed purple line shows the predicted efficiency if four-wave mixing is removed from the model.

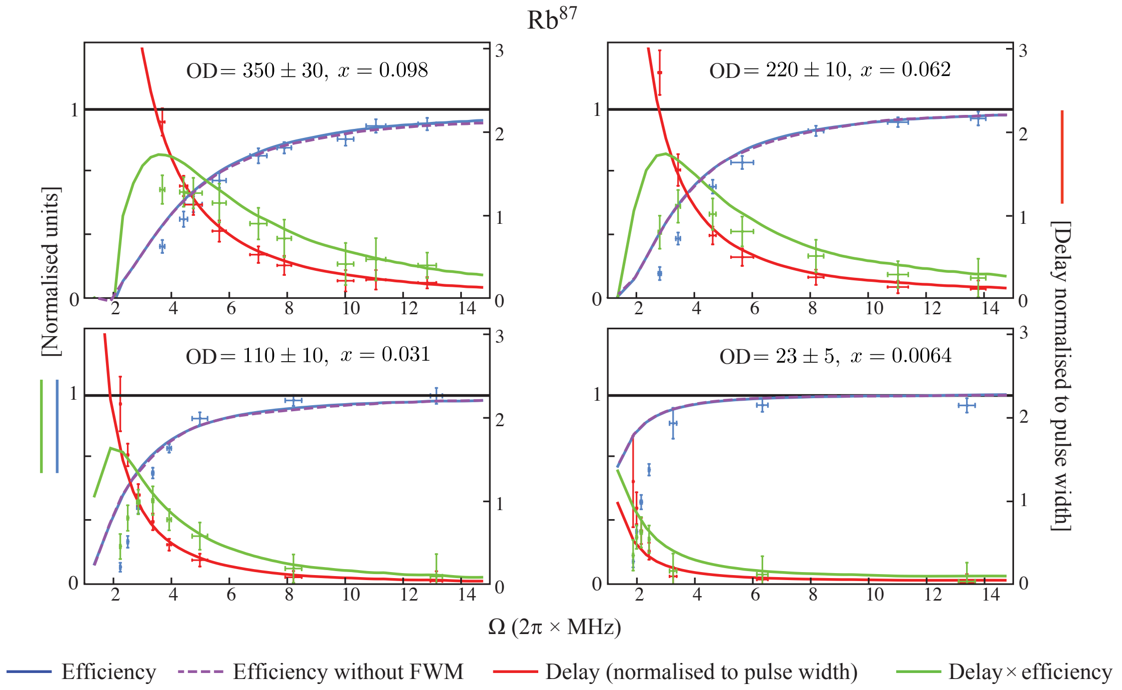

Figure 8.7: The efficiencies (blue) and delay times (red) for ^{85}Rb. The delay times are normalised to the 6 μs input pulse width. The product of delay and efficiency (green) is relative to the reference input. The solid lines show the predicted values based on eqs. (8.19). The dashed purple line shows the predicted efficiency if four-wave mixing is removed from the model.

discrepancy in observed 4WM is a result of the different effective 4WM strength in the experiment due to the lower optical depth and larger ground-state splitting in ^{87}Rb relative to ^{85}Rb. We verified that for ^{87}Rb high delays are indeed possible with minimal 4WM by repeating the experiment with that isotope. While the maximum achieved optical depth of $\mathrm{OD} = 350 \pm 30$, was lower than for ^{85}Rb, high delay-bandwidth products were still obtained, as is shown in fig. 8.7. Here the effective interaction strength ratio was calculated to have a theoretical value of $\eta_{\mathrm{eff}} = 1.33$, giving an effective 4WM strength of $x = 0.08$ for an optical depth of 350. At the resolution of the plot, a theoretical line that includes 4WM is indistinguishable from one that does not have 4WM. Our results are therefore consistent with the results of [a35].

The measured efficiencies fall below those predicted by the model for small optical depths and low control field powers. The effect is present in both species of Rb but is particularly pronounced in ^{87}Rb. This may be a result of the simplification to a four-level model as it does not include the effects of optical pumping due to the control field. Our analysis relies on the assumption that the conversion between the measured control field power and the resulting effective Rabi frequency remains constant across all of the data. This assumption is not valid if the distribution of population amongst the Zeeman sub-levels varies with the control field power or the optical depth. A more thorough analysis would include the effects of optical pumping, however, this is complicated by radiation trapping in optically dense ensembles [a70]. We further note that some departure from the predicted output may be a result of eddy currents in the optical bench that persist after the trapping magnetic fields are turned off.

The multi-modal capacity of EIT-based memory scales poorly with increasing optical depth; at best the modal capacity is $N \approx \sqrt{D}/3$ [a71]. Despite this, the optical depth achieved in our experiment was sufficient to demonstrate enough delay that two pulses are contained entirely within the ensemble simultaneously. Figure 8.8 shows results taken at an optical depth of (560 ± 40) demonstrating a delay-bandwidth product of ≈ 3.7, calculated by the ratio of the delay to transmitted pulse full-width-half-max at 50% efficiency. This is the first time that EIT has been observed at such large optical depths and the first demonstration of temporally multi-mode delay in EIT. The optical depth achieved for this data is slightly higher than that for the data presented in fig. 8.6. Furthermore, gain is clearly observed as we see regions where the transmission efficiency appears to be greater than unity.

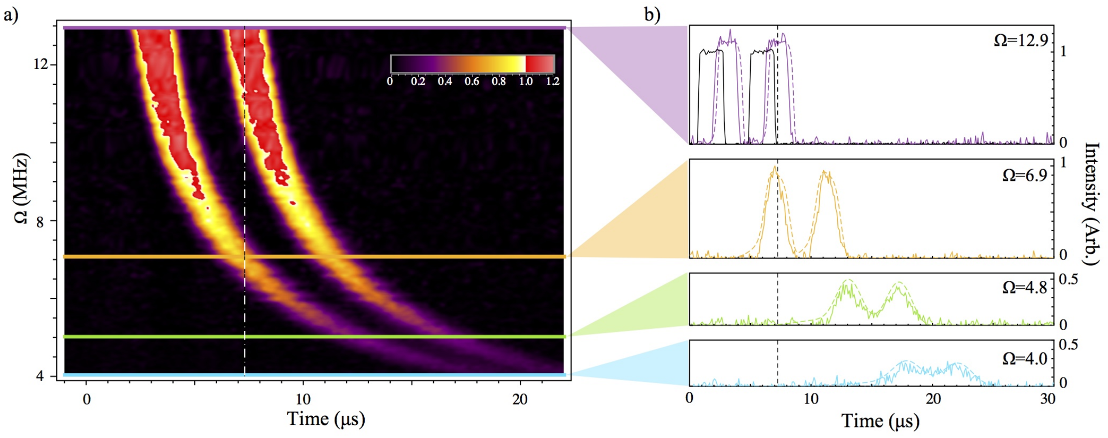

Figure 8.8: Temporally multimode delay of two signal pulses. Here, the large optical depth is used to slow the input pulses enough that they are both contained entirely within the atomic ensemble. Part (a) shows the ouput from the delay medium as the control field power is varied. Individual traces from this data are shown in (b) with the associated traces highlighted in the same color in part (a). The top trace of (b) also shows the input pulse as a reference, with some 4WM gain being apparent in the output.

CHAPTER 9

Gradient echo memory

A gradient echo memory (GEM) operates on the principle of time-reversal to recover an optical pulse. In this respect, GEM is not unique; a variety of photon-echo memory techniques based on this principle exist. The most common is the use of controlled reversible inhomogeneous broadening (CRIB) in an ensemble of atoms. This approach to recovering an optical states relies on reversing the direction of the phase evolution of each atom in an ensemble after a given time. As the atoms independently return to the state that they were in when the pulse was absorbed, the macroscopic polarisation of the medium is restored to same state as during the initial absorption. If re-absorption can be avoided, the optical pulse is recovered.

In a gradient echo memory, a narrow linewidth atomic transition is broadened by the application of an external field gradient along the direction of propagation. An optical pulse with a bandwidth that is matched to the broadened linewidth is then sent into the medium and is absorbed. The broadening of medium ensures that as the optical field drives the atomic coherence, the coherence de-phases and halts propagation of optical field, converting it entirely into a collective state of the atomic ensemble. Crucially, the de-phasing is coherent and controlled rather than the result of an incoherent process. Reversal of the gradient causes a re-phasing of the atomic coherence and thus a regeneration of the optical field. Re-absorption is avoided provided that the gradient is monotonic; each slice of the memory absorbs only one frequency, as the frequencies are regenerated and exit the medium, they pass only atoms that are not resonant.

In the three-level GEM scheme, a strong optical coupling field and a weak signal field form an off-resonant Raman transition in an ensemble of Λ-type three-level atoms. By applying a detuning gradient across the ensemble, a signal pulse can be transferred into a coherent atomic spin wave excitation between the two ground states and re-emitted when the gradient is reversed.

9.1 Theory

9.1.1 The GEM polariton

We start from the simplified equations of motion for a two-level atom, eqs. (6.44a) and (6.44b), using the normalized coordinates $\{\xi, \tau\}$. For the GEM protocol, a spatial gradient of the detuning is included such that

$$\partial_\tau \hat{P} = -\left(i\eta(\tau)\xi + \Gamma\right)\hat{P} + i\sqrt{d}\Gamma\hat{\mathcal{E}} \tag{9.1}$$

$$\partial_\xi \hat{\mathcal{E}} = i\sqrt{d}\hat{P}, \tag{9.2}$$

where $\eta(\tau)\xi$ is the spatial gradient, that can be varied in time and ξ now runs from $-1/2$ to $1/2$.

We can gain insight into the dynamics of the coupled atom-light interaction by examining eqs. (9.1) and (9.2) in the spatial Fourier domain. Both the optical field and atomic coherence are expanded in a spatial plane wave decomposition, $a(\xi) = \sum_k a_k e^{ik\xi}$, so that $\frac{\partial}{\partial\xi} \rightarrow ik$, $\xi \rightarrow -i\frac{\partial}{\partial k}$. Equations (9.1) and (9.2) can then be written

$$\partial_\tau \hat{P} = -\left(\eta\partial_k + \Gamma\right)\hat{P} + i\sqrt{d}\Gamma\hat{\mathcal{E}}, \tag{9.3}$$

$$k\hat{\mathcal{E}} = \sqrt{d}\hat{P}. \tag{9.4}$$

Taking a similar approach to the case of EIT, we can define

$$\hat{\Psi} = \cos\theta\hat{\mathcal{E}} + \sin\theta\hat{P} \tag{9.5}$$

$$\hat{\Phi} = \sin\theta\hat{\mathcal{E}} - \cos\theta\hat{P} \tag{9.6}$$

$$\tan\theta \equiv \frac{k}{\sqrt{d}} \tag{9.7}$$

where $\hat{\Phi} = 0$. Noting that $\hat{\mathcal{E}} = \cos\theta$, $\hat{P} = \sin\theta$, and $\partial_k\theta = \cos\theta/\gamma$ we find the equation of motion of $\hat{\Psi}$ to be

$$\left(\partial_\tau + \eta\partial_k + \Gamma + \left(\frac{\eta}{k^2 + d} - i\Gamma\right)\frac{d}{k}\right)\hat{\Psi} = 0. \tag{9.8}$$

We can examine each term of the propagation equation to determine the behaviour of the polariton during the storage and recall process. The first two terms $\partial_\tau + \eta\partial_k$ describe a shape-preserving advection in k-space for the GEM polariton, much as eq. (8.17a) describes an advection in space for the EIT polariton. Γ is a constant loss rate due to de-phasing, while the remaining term can be broken down into two parts. The $i\Gamma d/k$ term introduces an accumulation of phase in k-space near $k = 0$. This corresponds to the forward propagation of the polariton in real space when the polariton re-phases. The non-unitary term $\eta d/((k^2 + d)k)$ is negligible when $d\Gamma/\eta \gg 1$, corresponding to the situation where the effective optical depth of the broadened ensemble is much larger than 1. If this condition is not satisfied, the optical field is not entirely converted into the polariton, hence a non-unitary evolution.

Equations (9.5), (9.7) and (9.8) have an intuitive consequence: the propagation of the pulse in real space and the magnitude of the electric field are linked. If the optical field provides a significant contribution to the polariton, it will propagate in real space; therefore, storage of the optical field requires that the polariton evolve to a sufficiently high k value that the polariton is entirely atomic. This can be interpreted a requirement that the applied broadening of the excited state be larger than the bandwidth of the input pulse. The polariton picture provides an intuitive approach to the memory dynamics and leads to the observation that the storage and recall dynamics can be controlled through the gradient $\eta(t)$. Figure 9.1 shows the memory operation in both real space and k-space to illustrate the dynamics.

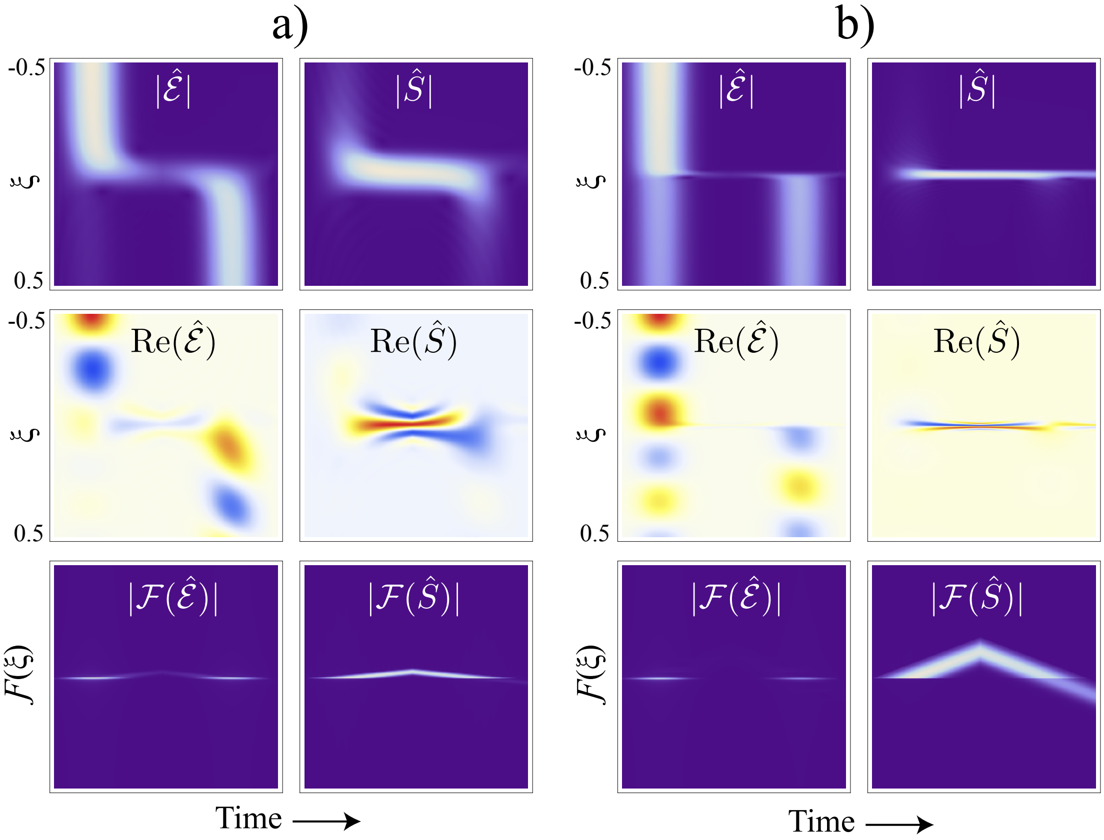

Figure 9.1: Simulation of a two-level GEM protocol for two values of the gradient. The simulations solve eqs. (6.66) with $\Gamma = \pi \times 6.07$ MHz, $\Delta = 2\pi \times 200$ MHz, $\Omega = 2\pi \times 3$ MHz, $\gamma = \pi \times 1$ kHz, and $d = 1000$. The FWHM of the input is 10 μs and the storage time is 40 μs. The gradients span $2\pi \times 0.8$ MHz in (a) and $2\pi \times 3.2$ MHz in (b). For each simulation, the absolute values and real parts of $\hat{\mathcal{E}}$ and $\hat{S}$ are shown to illustrate how the coherence propagates in terms of magnitude and phase. For the smaller gradient, a pulse chirp can be seen in the real part of the recalled light due to an imperfect time-reversal symmetry. The magnitudes of the fields are also shown in k-space to illustrate the propagation of the polariton.

The polariton can be stored at a high k value, that is, the spin-wave is dephased, and be brought back to $k = 0$ for recall when required. This provides the memory protocol with an on-demand character: The gradient field need not be switched half way between storage and recall, but can but turned off for an arbitrary wait time and then reversed when a recall is needed. This control also provides multi-mode capacity by allowing multiple input pulses to be stored at various k values and then recalled in reverse order.

9.1.2 Transfer functions

While the polariton picture is a useful tool for intuitively examining the memory dynamics, it is also useful to examine the equations that govern the read and write phases of memory operation more directly. We can analyse the equations of motion analytically by following the approach laid out in [a72, a73]. We assume that η is constant in time and that $\hat{P}(\tau = 0) = 0$ to solve the equations of motion for the write phase of the memory. Taking a

Laplace transform of eq. (9.1) yields

$$\hat{P}(s) = \frac{i\sqrt{d}\Gamma}{s+\Gamma+i\eta\xi}\hat{\mathcal{E}} \tag{9.9}$$

which we can insert into eq. (9.2) and integrate to obtain

$$\hat{\mathcal{E}}(\xi,s) = \left(\frac{s+\Gamma+i\eta\xi}{s+\Gamma+i\eta\xi_0}\right)^{id\frac{\Gamma}{\eta}} \cdot \hat{\mathcal{E}}(\xi_0,s), \tag{9.10}$$

which we can write in terms of the optical fields at the input and output of the ensemble. The output optical field can be found explicitly in the time-domain as a convolution as

$$\hat{\mathcal{E}}_{out}(\tau) = -\Gamma d\int_0^\tau d\tau' \left(e^{-(\Gamma+i\eta/2)(\tau-\tau')}{}_1F_1(1+id\Gamma/\eta,2,i\eta(\tau-\tau'))\right)\hat{\mathcal{E}}_{in}(\tau') \tag{9.11}$$

where ${}_1F_1$ is a confluent hypergeometric function. We can also find the form of the atomic coherence that is left in the ensemble after writing into the memory

$$\hat{P}(\xi,s) = \frac{i\sqrt{d}\Gamma}{s+\Gamma+i\eta\xi}\left(\frac{s+\Gamma+i\eta\xi}{s+\Gamma+i\eta\xi_0}\right)^{id\frac{\Gamma}{\eta}} \cdot \hat{\mathcal{E}}(\xi_0,s) \tag{9.12}$$

which can also be transformed into the time domain by taking a convolution

$$\hat{P}(\xi,\tau) = i\sqrt{d}\int_0^\tau d\tau' \left(e^{-(\Gamma+i\eta\xi)(\tau-\tau')}L(-id\Gamma/\eta, i(\tau-\tau')(1/2+\xi)\eta)\right)\hat{\mathcal{E}}_{in}(\tau') \tag{9.13}$$

where L is a Laguerre polynomial.

These equations govern the dynamics of the memory during the write phase of the memory. The read-phase dynamics can be found by using the same method with the boundary conditions $\hat{P}(\xi,\tau=0)=\hat{P}_{in}(\xi)$ and $\hat{\mathcal{E}}_{in}=0$. Again taking a Laplace transform of the equations of motion, we arrive at

$$\hat{\mathcal{E}}(\xi,s) = i\sqrt{d}\int_{-1/2}^{\xi} \frac{d\xi'}{s+\Gamma+i\eta\xi'}\left(\frac{s+\Gamma+i\eta\xi}{s+\Gamma+i\eta\xi'}\right)^{id\frac{\Gamma}{\eta}} \cdot \hat{P}_{in}(\xi') \tag{9.14}$$

$$\hat{\mathcal{E}}_{out}(\tau) = i\sqrt{d}\int_{-1/2}^{1/2} d\xi' e^{-(i\eta\xi'+\Gamma)\tau}L(id\Gamma/\eta,(i\eta(\xi'-1/2))\tau) \cdot \hat{P}_{in}(\xi') \tag{9.15}$$

and

$$\hat{P}(\xi, s) = \frac{1}{s + i\eta\xi + \Gamma} \int_{-1/2}^{1/2} d\xi' \left(\delta(\xi - \xi') - \frac{d\Gamma}{s + i\eta\xi' + \Gamma} \left(\frac{s + \Gamma + i\eta\xi}{s + \Gamma + i\eta\xi'} \right)^{id\frac{\Gamma}{\eta}} \right) \cdot \hat{P}_{in}(\xi_0, s) \tag{9.16}$$

$$\hat{P}_{out}(\xi, \tau) = e^{-(i\eta\xi + \Gamma)\tau} \int_{-1/2}^{\xi} d\xi' \left(\delta(\xi - \xi') - d\Gamma\tau {}_1F_1(1 + id\Gamma/\eta, 2, i\eta\tau(\xi - \xi') \right) \cdot \hat{P}_{in}(\xi_0, s). \tag{9.17}$$

Equations (9.11) and (9.13) don't really provide an intuitive picture of the memory dynamics. We can however gain some insight from the transfer functions that relate the input optical field to the stored spin coherence and the optical field that doesn't get absorbed. Defining the write-phase transfer functions according to

$$\hat{\mathcal{E}}_{out}^{write}(\tau) = \int_{0}^{\tau} g_{write}(\tau - \tau')\hat{\mathcal{E}}_{in}(\tau')d\tau' \tag{9.18}$$

$$\hat{P}_{out}^{write}(\xi) = \int_{0}^{T} h_{write}(\xi, T - \tau')\hat{\mathcal{E}}_{in}(\tau')d\tau' \tag{9.19}$$

where T is a characteristic write time for the memory. We can further write the transfer functions for the read phase of the memory

$$\hat{\mathcal{E}}_{out}^{read}(\tau) = \int_{-1/2}^{1/2} h_{read}(\xi', \tau)\hat{P}_{in}(\xi')d\xi' \tag{9.20}$$

$$\hat{P}_{out}^{read}(\xi) = \int_{-1/2}^{\xi} g_{read}(\xi, \xi', T)\hat{P}_{in}(\xi')d\xi'. \tag{9.21}$$

With these definitions, we can write down the transfer functions in the Laplace domain as

$$\tilde{g}_{write}(s) = \left(\frac{s + i\eta/2 + \Gamma}{s - i\eta/2 + \Gamma} \right)^{id\frac{\Gamma}{\eta}} \tag{9.22}$$

$$\tilde{h}_{write}(s, \xi) = \frac{i\sqrt{d}\Gamma}{s + i\eta\xi + \Gamma} \left(\frac{s + i\eta\xi + \Gamma}{s - i\eta/2 + \Gamma} \right)^{id\frac{\Gamma}{\eta}} \tag{9.23}$$

$$\tilde{g}_{read}(\xi, \xi', s) = \frac{1}{s + i\eta\xi + \Gamma} \left(\delta(\xi - \xi') - \frac{d\Gamma}{s - i\eta\xi' + \Gamma} \left(\frac{s + i\eta\xi + \Gamma}{s - i\eta\xi' + \Gamma} \right)^{id\frac{\Gamma}{\eta}} \right) \tag{9.24}$$

$$\tilde{h}_{read}(s, \xi) = \frac{i\sqrt{d}}{s + i\eta\xi + \Gamma} \left(\frac{s + i\eta/2 + \Gamma}{s + i\eta\xi + \Gamma} \right)^{id\frac{\Gamma}{\eta}}. \tag{9.25}$$

The functions $\tilde{h}_{write}(s, \xi)$ and $\tilde{h}_{read}(s, \xi)$ that respectively describe the mappings $\hat{\mathcal{E}} \to \hat{P}$ and $\hat{P} \to \hat{\mathcal{E}}$ show a symmetry that underlies the memory process. Some intuition can be gained by plotting the transfer functions in the Fourier domain by setting $s = i\omega$. A plot of the absolute value and phase of $\tilde{g}_{write}(i\omega)$, fig. 9.2(a,b), shows the frequency response of the ensemble for $\hat{\mathcal{E}}_{in}$. The absorption window that defines the memory bandwidth can easily be seen. We can find the absorption depth of the broadened ensemble by calculating the value of $\tilde{g}_{write}(i\omega)$ for $\omega \ll \eta$. Under this assumption, the effective optical depth is

$$\tilde{g}_{write} \approx \exp\left(-2d\frac{\Gamma}{\eta}\arctan\left(\frac{\eta}{2\Gamma}\right)\right) \tag{9.26}$$

which for $\eta \gg \Gamma$ can be further approximated as $\exp\left(-\pi d\Gamma/\eta\right)$. This is shown on fig. 9.2 as the dashed line.

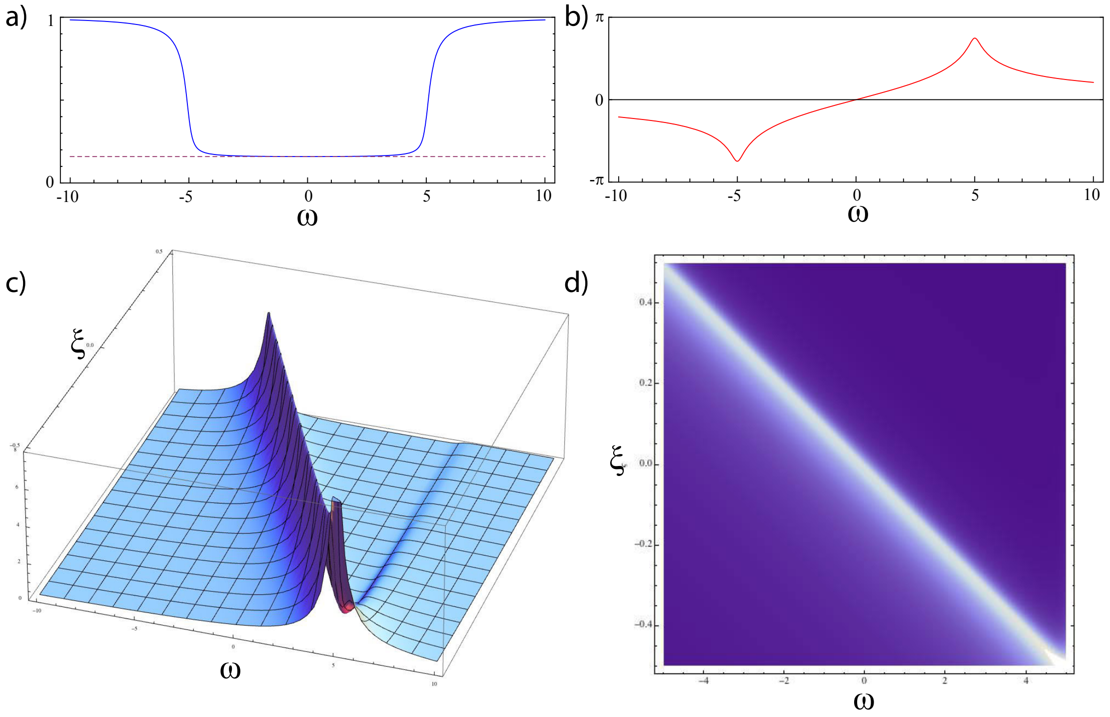

Figure 9.2: a,b) The absolute value (blue) and phase (red) of the transfer function $\tilde{g}_{write}(i\omega)$ with the parameters $\Gamma = 0.2$, $\eta = 10$, $d = 6$. A small frequency approximation of the absolute value is shown as the purple dashed trace. c,d) The absolute value of the transfer function $\tilde{h}_{write}(i\omega, \xi)$ with the parameters $\Gamma = 0.2$, $\eta = 10$, $d = 6$.

Plotting $\tilde{h}_{write}(i\omega, \xi)$, shown in fig. 9.2(c,d) illustrates the nature of the storage process by showing the linear mapping between the Fourier components of the input pulse and spatial location within the memory. We can also see this behaviour by examining the read and write transfer functions

$$\begin{aligned} h_{write}(\tau, \xi) &= \mathcal{L}^{-1}\{\tilde{h}_{write}(s, \xi)\}(\tau) && (9.27)\\ &= i\sqrt{d}\Gamma e^{-\Gamma\tau} e^{-i\eta\xi\tau} L(-id\Gamma/\eta, i\eta\tau(1/2+\xi)) && (9.28)\\ &\approx i\sqrt{d}\Gamma e^{-i\eta\xi\tau} e^{-id\frac{\Gamma}{\eta}\ln(\eta\tau(\xi+1/2))} e^{-\Gamma\tau} \frac{e^{-d\frac{\pi}{2}\frac{\Gamma}{\eta}}}{\tilde{\Gamma}\left(1-id\frac{\Gamma}{\eta}\right)} && (9.29) \end{aligned}$$

and

$$\begin{aligned}
h_{read}(\tau,\xi) &= \mathcal{L}^{-1}\{\tilde{h}_{read}(s,\xi)\}(\tau) && (9.30)\\
&= i\sqrt{d}e^{-\Gamma\tau}e^{-i\eta\xi\tau}L(id\Gamma/\eta, i\eta\tau(1/2+\xi)) && (9.31)\\
&\approx i\sqrt{d}e^{-i\eta\xi\tau}e^{-id\frac{\Gamma}{\eta}\ln(\eta\tau(\xi+1/2))}e^{-\Gamma\tau}\frac{e^{d\frac{\pi}{2}\frac{\Gamma}{\eta}}}{\tilde{\Gamma}\left(1+id\frac{\Gamma}{\eta}\right)} && (9.32)
\end{aligned}$$

where $\tilde{\Gamma}$ is the Gamma function, not to be confused with the excited state decay rate Γ. The first exponential term acts as a Fourier kernel between τ and $\tilde{\xi} = \eta\xi$ under the convolution integrals. This transform, however, is not ideal due to the logarithmic phase added by the following term. This phase term grows slowly compared to the Fourier kernel in the limit of $\eta\xi\tau \gg 1$, which corresponds to long storage times. The next term is due to the decay of the excited state and degrades memory performance for storage times that are not short relative to Γ. The final term provides the overall efficiency of the process and can be removed from the convolution integrals. Defining

$$\xi_{read/write} \equiv \frac{e^{\pm d\frac{\pi}{2}\frac{\Gamma}{\eta}}}{\tilde{\Gamma}\left(\pm id\frac{\Gamma}{\eta}\right)}, \tag{9.33}$$

the efficiency of the memory process is given by $|\xi_{read/write}|^2$. By applying some functional relations of the Gamma function, we arrive at a simple expression for the efficiency

$$|\xi_{read/write}|^2 = \frac{1}{\pi}\frac{\eta}{\Gamma}\left(1-e^{-2\pi d\frac{\Gamma}{\eta}}\right). \tag{9.34}$$

This expression shows that the fraction of the excitation that is transferred to the atoms is $1-\exp\left(-2\pi d\Gamma/\eta\right)$, while the factor $\eta/(\pi\Gamma)$ reflects the fact that the excitation is spread out across the ensemble. The memory operation that includes both the write and read steps will involve successive applications of h_{write} and h_{read}. This will transform optical information into an atomic collective excitation that maps the Fourier spectrum of the input pulse into spatial location in the memory and later reverses the process to restore the optical pulse. This will result in a total amplitude efficiency of

$$\pi\frac{\Gamma}{\eta}|\xi_{write}\times\xi_{read}| = 1-e^{-2\pi d\frac{\Gamma}{\eta}} \tag{9.35}$$

9.1.3 GEM as a beamsplitter

While the detailed dynamics of the memory are complicated, the magnitude of the transfer functions under the conditions associated with memory operation can be expressed as simple beam-splitter relations with the transmission T and reflection R defined by

$$\begin{aligned}
T &= e^{-2\beta} && (9.36a)\\
R &= 1-e^{-2\beta} && (9.36b)\\
\beta &\equiv \pi d\frac{\Gamma}{\eta}. && (9.36c)
\end{aligned}$$

This beam-splitter acts between an optical mode propagating in space and a stationary atomic collective excitation that propagates in time. The exact temporal (spectral) shape of the optical pulse and the spatial shape of the excitation are related according to the transfer functions $h_{read/write}$. For storage times that are long compared to the pulse durations, $h_{read/write}$ is approximately a Fourier transform and will preserve the shape of Gaussian input pulses.

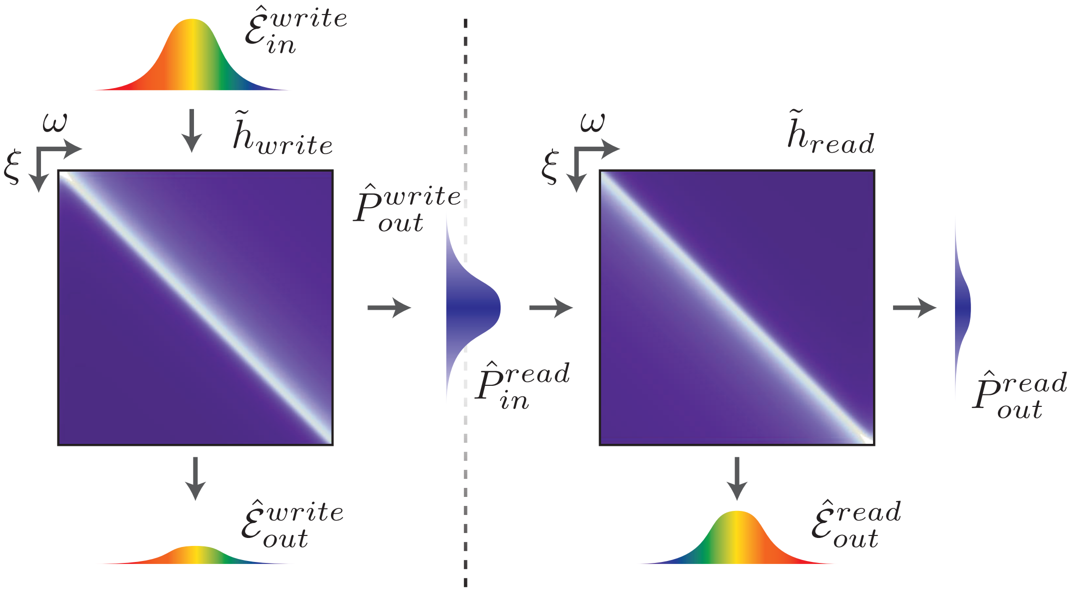

Figure 9.3: Conceptual model of GEM as a beam-splitter.

The concept of the GEM protocol acting as beam-splitter provides an intuitive picture of how an excitation is preserved through interaction with the memory and offers a simplified model inefficiency in the memory. This simple picture also forms the basis of a number of ideas that can be used to extend the utility of the memory to perform much more complex operations than simple storage and retrieval. These will be explored in the next chapters and form the core component of the original work presented in this thesis.

9.2 Three-level GEM

The gradient echo memory concept can be immediately applied to three-level atomic systems that are driven off-resonance by noting that, in the far detuned regime, they obey the same dynamics as two-level systems. Referring to eqs. (6.69), we can directly apply the results of sections 9.1.1 and 9.1.2 to the three-level case but with direct optical control of the effective atomic coherence through the amplitude of the control field. Explicitly, the quantities for the effective two-level atom that approximates the Raman absorption line are

$$\Delta' = \delta - \Delta\kappa \tag{9.37a}$$

$$\Gamma' = \gamma + \Gamma\kappa \tag{9.37b}$$

$$\hat{P}' = \frac{\Omega}{\Delta}\hat{S} \tag{9.37c}$$

$$\kappa \equiv \frac{|\Omega|^2}{\Gamma^2 + \Delta^2} \approx \frac{|\Omega|^2}{\Delta^2}. \tag{9.37d}$$

The term $\Delta\kappa$ is the ac-Stark shift, which must be accounted for in the two-photon detuning, $\Gamma\kappa$ is the linewidth of the effective of atom, including both the natural decay rate of $\hat{S}$ and the power-broadening due to the control field and Ω/Δ is a rescaling of $\hat{S}$ that can be optically controlled.

An estimation of the efficiency of a three-level gradient echo memory can be made from these effective parameters. The three sources of loss to consider are: transmission through the ensemble because the effective optical depth is too low, loss due to decoherence of $\hat{S}$ as a result of power broadening, and off-resonant absorption into the excited state. Using the notation λ for the efficiency, the transmission due to insufficient optical depth is given by eq. (9.36b) as

$$\lambda_d = \sqrt{1 - e^{-4\pi d\kappa\frac{\Gamma}{\eta}}} \approx 1 - \frac{1}{2}e^{-4\pi d\kappa\frac{\Gamma}{\eta}}. \tag{9.38}$$

If we assume that the control field is left off during storage and neglect the intrinsic decay rate of $\hat{S}$, we can approximate the loss due to power broadening as the scattering rate $\Gamma\kappa$ times the pulse width. Since η must roughly match the pulse bandwidth, the efficiency is

$$\lambda_\Omega = e^{-\kappa\frac{\Gamma}{\eta}}. \tag{9.39}$$

Finally, the efficiency limited by off-resonant absorption is

$$\lambda_\Gamma = e^{-d\frac{\Gamma^2}{\Delta^2}}. \tag{9.40}$$

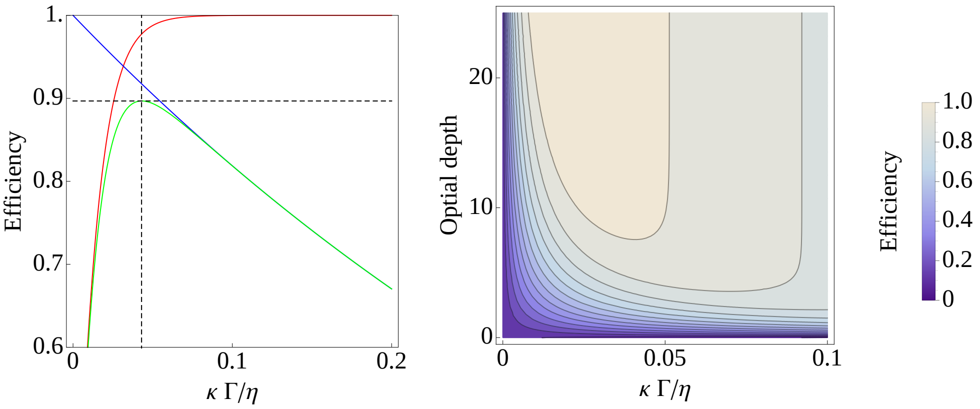

Figure 9.4: The estimated efficiency λ^2 as a function of a characteristic parameter $\kappa\Gamma/\eta$. The left plot shows the efficiency λ^2 in green with the contributions λ_d^2 and λ_Ω^2 in red and blue, respectively. The approximated maximum efficiency is shown by the dashed lines. The optical depth used for the plot was 7.3. The right plot shows Λ^2 as a function of both the characteristic parameter and the optical depth d.

With sufficient control field power available, Δ can be made large enough to render the off-resonant absorption negligible. In this limit, the total loss will be

$$\lambda = \lambda_d \cdot \lambda_\Omega \approx e^{-\kappa\Gamma/\eta} - \frac{1}{2}e^{-(1+4\pi d)\kappa\Gamma/\eta}. \tag{9.41}$$

From this we can find that, when $4\pi d \gg 1$, the efficiency reaches a maximum value of

$$\lambda \approx \left(\frac{1}{2\pi d}\right)^{\frac{1}{4\pi d}} \left(1 - \frac{1}{4\pi d}\right) \tag{9.42}$$

when the control field is set such that

$$\kappa\frac{\Gamma}{\eta} \approx \frac{\log(2\pi d)}{4\pi d}. \tag{9.43}$$

The estimated efficiency λ^2 is shown in fig. 9.4.

9.3 A high-efficiency memory

Predictions of memory efficiency and general performance must, of course, be verified by experiments. In the warm-atom memory experiment efficiencies of up to 87% were obtained by careful optimisation of the memory parameters [a3]. In the cold-atom memory, efficiencies of 80% were achieved [a5]. In the case of the cold-atom memory, the results can be readily compared with a simple theoretical model thanks to the negligible inhomogeneous broadening of the excited state. Figure 9.5 shows experimental results for a storage efficiency of 79% compared with a model using the same parameters. The optical depth used in the simulation, 347, was determined by measurements of the unbroadened Raman lines. This is roughly in agreement with the measured resonant optical depth of ≈ 500.

Because only the atoms in one of the Zeeman sub-levels are used for the memory a circularly polarised optical pumping beam was applied to prepare the ensemble into a single sub-level. The optical depth inferred from the Raman lines suggests that this pumping was approximately 70% effective. Using the other experimental parameters in the model, excellent agreement is found with the experiment.

9.4 The quantum performance of GEM

A candidate quantum memory must be simultaneously efficient and add little noise to the recalled state. Noise could result from technical sources in the experiment, such as phase-noise from acoustic vibrations, or could be intrinsic to the memory, such as the addition of noise photons from a non-linear process such as four-wave-mixing. This latter source is particularly concerning due to the evidence of 4WM seen in the EIT experiment. Before exploring various possibilities for using the gradient echo memory, it is important to present evidence that the memory can provide a high quantum fidelity.

We performed the characterisation using a hot vapour memory, as described in section 7.3. This memory had already been demonstrated to operate with an efficiency of 87%.

9.4.1 Homodyne detection

In order to characterise the quantum performance of the memory, a detector that is sensitive to the level of single photons must be used. One possibility is to use single-photon counting detectors that rely on an avalanche cascade of electrons in the sensitive semiconductor to amplify single excitations into large signals. There are two downsides to this approach. One is that phase information is not recorded, making it less straightforward

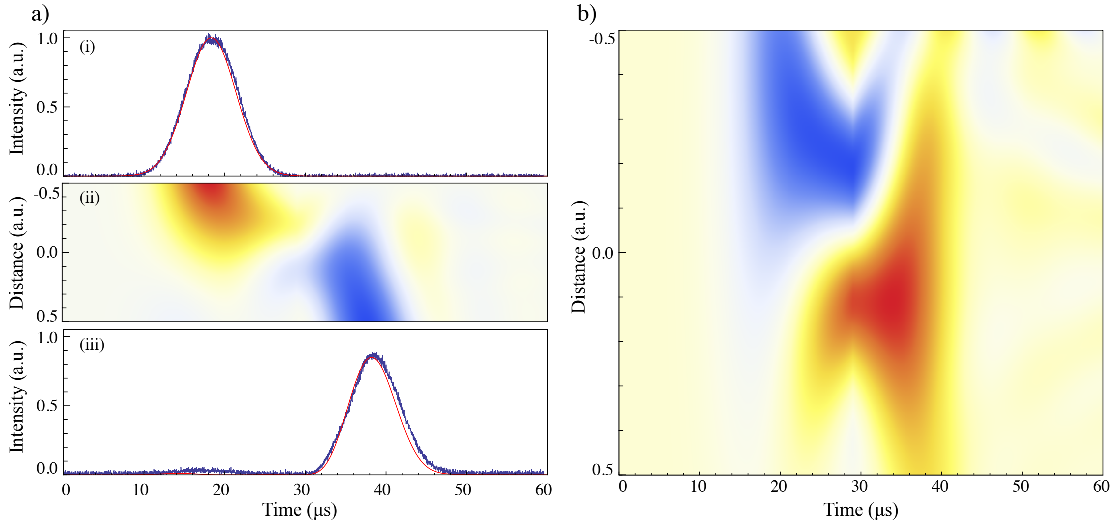

Figure 9.5: Experimental results of a cold-atom gradient echo memory shown along with numerical simulations that illustrate the memory dynamics. a) The optical field and the input (i), in memory (ii) and at the output (iii). (i) and (iii) correspond to intensity while (ii) shows the real part of the amplitude. b) A numerical simulation of $\hat{S}$ during storage. In real experiments the memory evolves to low k values as higher gradients reduce the effective optical depth of the ensemble. The simulations solve eqs. (6.66) with $\Gamma = \pi \times 6.07$ MHz, $\Delta = 2\pi \times 200$ MHz, $\Omega = 2\pi \times 3$ MHz, $\gamma = \pi \times 1$ kHz, and $d = 347$. The gradients give a broadening of 143 kHz for the input and 175 kHz for the output. The asymmetry of the gradient compensates for pulse distortion due to the breaking of time-reversal symmetry in the low k-value memory.

to demonstrate that the phase is preserved through the memory. The second is that the bright control fields necessary for GEM are difficult to filter out to below the single photon level. The alternative is to employ homodyne detection.

In a homodyne detector, the weak quantum field is mixed with a bright reference coherent on a beam-splitter. The interference produced between the quantum field and the coherent state is proportional to the amplitude of each field, allowing the magnitude of the quantum signal to be amplified to a level that high-efficiency linear photo-diodes can resolve. Fluctuations in the bright field are removed by ensuring that the beam-splitter balances each of the two input modes equally into the two output modes. If both output modes are detected and the resultant photo-currents are subtracted, only the beat note between the the two fields remains.

Explicitly, the modes emerging from a 50/50 beam-splitter, $\hat{a}_1'$ and $\hat{a}_2'$, can be written in terms of the input modes, $\hat{a}_1$ and $\hat{a}_2$ [b74]

$$\hat{a}_1' = \frac{1}{\sqrt{2}}(\hat{a}_1 - \hat{a}_2), \tag{9.44}$$

$$\hat{a}_2' = \frac{1}{\sqrt{2}}(\hat{a}_1 + \hat{a}_2). \tag{9.45}$$

The number of photons arriving at each detector is

$$\hat{n} = \hat{a}'^{\dagger}\hat{a}' \tag{9.46}$$

$$= \frac{1}{2}(\hat{a}_1 \pm \hat{a}_2)^{\dagger}(\hat{a}_1 \mp \hat{a}_2) \tag{9.47}$$

$$= \frac{1}{2}(\hat{n}_1 + \hat{n}_2 \mp \hat{a}_1^{\dagger}\hat{a}_2 \pm \hat{a}_2^{\dagger}\hat{a}_1). \tag{9.48}$$

Subtracting the photon numbers arriving at each detector then gives

$$\hat{n}_{21} = \hat{n}_2' - \hat{n}_1' = \hat{a}_1^{\dagger}\hat{a}_2 - \hat{a}_2^{\dagger}\hat{a}_1. \tag{9.49}$$

Treating the local oscillator as a bright, classical field the photon number difference is then

$$\hat{n}_{21} = \alpha_{LO}^{*}\hat{a}_s + \alpha_{LO}\hat{a}_s^{\dagger} \tag{9.50}$$

$$= \sqrt{(2)}|\alpha_{LO}|\hat{q}_{\theta} \tag{9.51}$$

where α_{LO} is the amplitude of the local oscillator and $\hat{q}_{\theta}$ is the projection of $\hat{a}$ onto the phase θ defined by the local oscillator. The terms associated with either the local oscillator or the signal alone have been subtracted out, meaning that technical noise in the local oscillator does not appear in the signal.

The signal obtained by subtracting the photo-currents of the two detectors is thus proportional to a projection of $\hat{a}$ onto the bright local oscillator. By varying the phase of the local oscillator, we can measure the projection of $\hat{a}$ onto any angle in the complex plane. With a sufficient number of these projective measurements, the quantum state of the signal can be reconstructed.

In the experiment, we use two separate photo-detectors each with two stages of gain to provide low electronic noise and a large dynamic range. The analogue electronic signals from the photo-detectors are subtracted by combining them in an RF power-splitter with an adjustable phase phase delay added to one of the signals. The subtracted signal is then digitised by an National Instruments data acquisition card and stored.

State tomography

To measure the noise performance of the memory and accomplish quantum state tomography we recorded more than 10^5 homodyne measurements for each input and output state. The input pulses have a duration of 2 μs and were stored for 3 μs. The control field was switched off for 1 μs during storage to minimise decoherence due to scattering. We integrate the amplitude of the pulse over the pulse duration and associate one quadrature value for each pulse.

Figure 9.6 illustrates the scheme used to analyse the memory. The output from the memory is mixed on a 50/50 beam-splitter with a bright local oscillator and the photo-currents associated with each output from the beam-splitter are detected, subtracted and recorded, as shown in fig. 9.6 (a). The local oscillator is spatially mode-matched to the signal to ensure the highest possible interference contrast on the beam-splitter, and thus a high detection efficiency.

Because the measurement is pulsed, is was difficult to impose a fixed phase between the local oscillator and the signal using a locking technique. In fact, the phase drift between these was quite significant due to the large path length of the interferometer and because

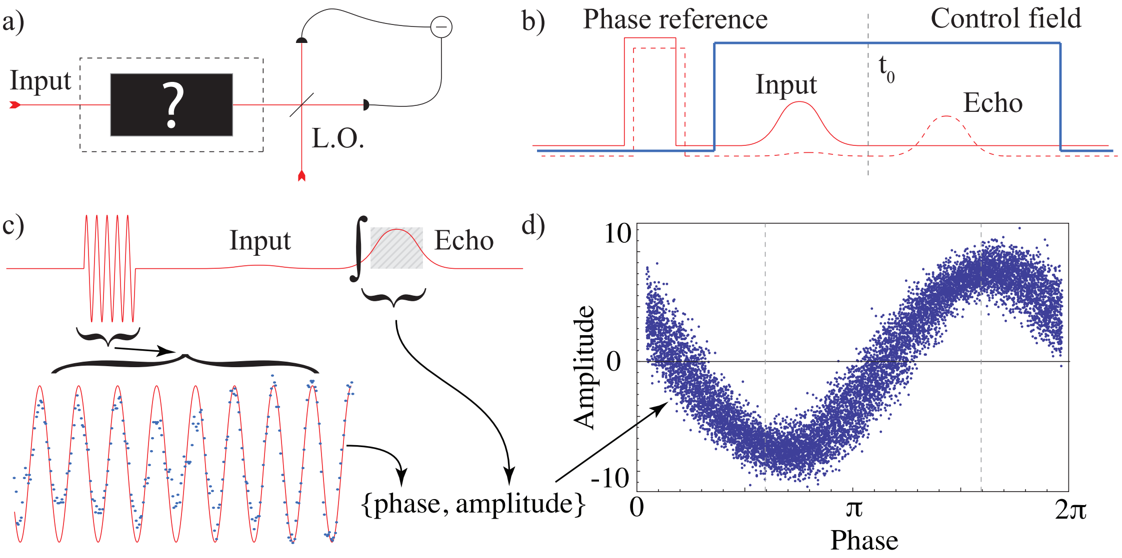

Figure 9.6: a) The conceptual layout for homodyne tomography of stored states. An input state is injected into the memory which has an unknown effect on the state. After the memory, the state is mixed with a local oscillator on a 50:50 beam-splitter. The two outputs of the beam-splitter are detected by a balanced detector that subtracts the photo-currents. b) The timing sequence. A bright phase reference pulse is sent though the memory prior to the memory operation. This pulse is at a different frequency to the probe. Next, the control field is turned on and an input pulse is sent into the memory. After the magnetic field gradient is reversed at t_0 an echo is recalled. Each homodyne data point is derived from a trace such as is shown in (c). The relative phase difference between the signal field and local oscillator is determined by fitting a sine function to the phase reference, with a sample fit shown in the inset. The amplitude of the associated quadrature in the echo is found by integrating the echo signal. The amplitude of the vacuum is obtained by integrating an equally sized window from a portion of the trace that contains no input light (not shown). All of the data points are combined to form a set of projective measurements of the signal field (d).

one arm contained the memory, a second Rb gas cell used for filtering and the optics needed to pulse the signal and expand the beam to a large diameter for propagation through the vapour cells. To determine the phase of a given homodyne detection, we used a bright phase reference sent through the same mode as the signal field but with the control field turned off. This pulse was transmitted through memory just prior to each storage and recall event and had a frequency offset relative to the signal and local oscillator so that an interference fringe could be observed. Figure 9.6 (b) shows the timing of the pulses.

For each homodyne measurement, the phase of the bright reference was fitted and the intensity of the homodyne signal during the echo was integrated. Figure 9.6 (c) shows how each trace of the homodyne signal was analysed and turned into one quadrature value. The process was repeated for inputs with the control fields left off, to provide a reference for the signals when no memory is present. This analysis of these differs only in the timing of the integration window. Additionally, a segment of each trace that contains no signal is integrated with the same window as for the signal to provide a point of shot noise.

All of the homodyne measurements are then combined to form one dataset that can be used to analyse the quantum states of the signal and input. The quadrature data are normalised such that the variance of the shot noise is unity. Figure 9.6 (d) shows a

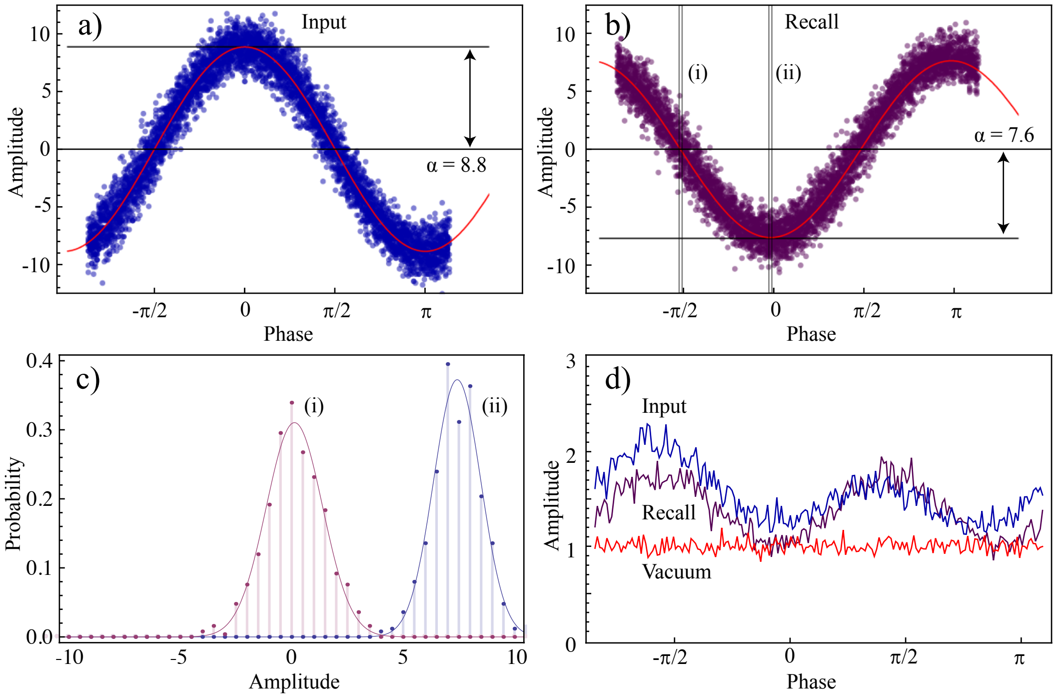

Figure 9.7: Homodyne data for a coherent state before (a) and after (b) storage in the gradient echo memory. Quadrature variances are calculated by binning the data into 200 segments of 500 samples. (c) shows histograms of the the segments (i) and (ii) shown in (b) with the corresponding normal distribution shown as a solid line. The calculated variances for the input and out states are shown in (d) relative to the vacuum noise.

sample dataset of homodyne measurements (or rather a subset of 10^4 points) for a stored and recalled pulse. The mean photon number of the input pulse, $\overline{n} = |\alpha|^2$ in this case is estimated to be 77.

A comparison of homodyne data taken for input and output states is shown is fig. 9.7. The mean photon number reduces from $\overline{n} = 77$ at the input (a) to $\overline{n} = 58$ after storage (b), indicating a storage efficiency of 75%. The data is binned into subsets of 500 points to determine the variance as a function of the projection angle. Histograms of the measurement values are shown in fig. 9.7 (c) for the phase (i) and amplitude (ii) quadratures. Plotting all of the variances (d) shows that there is phase noise in both the input and output. The source of this noise is likely from a combination of two sources: one is the accuracy of the phase fitting to find the relative local oscillator phase for each trace, the other is phase noise on the laser itself. The fitting uncertainty was reduced by tailoring the amplitude and frequency of the reference pulse. The technical noise from the laser is a result of a 87 kHz modulation used internally in the laser to lock an intra-cavity filtering etalon. This source was reduced for later experiments by purchasing a Ti:sapphire laser from a manufacturer that guarantees a lower modulation depth. The homodyne experiment, however, was not repeated with this new light source.

In [a49], we present various representations of the quantum state that has been recalled from the memory along with the input state. We can use these to quantitatively assess the effect of storage on the input states using measures such as the fidelity between the input and output states. These representations first require the photon-number density matrix of the optical state to be reconstructed from the homodyne data. This can be done by an iterative maximum-likelihood algorithm to find the density matrix that corresponds to a probability density function for the quadrature measurements that maximises the likelihood function for the given dataset. The algorithm [a75] uses successive guesses of the density matrix given by

$$\hat{\rho}^{(k+1)} = \mathbf{N}\left[\hat{R}(\hat{\rho}^{(k)})\hat{\rho}^{(k)}\hat{R}(\hat{\rho}^{(k)})\right], \tag{9.52}$$

where

$$\hat{R}(\hat{\rho}) = \sum_i \frac{\hat{\Pi}(\theta_i, x_i)}{pr_{\theta_i}(x_i)}, \tag{9.53}$$

$$\Pi_i = \sum_{n,m} \sqrt{\frac{2}{\pi}} e^{i\theta(n-m)} \frac{H_n(\sqrt{2}x)H_m(\sqrt{2}x)}{\sqrt{2^{n+m}n!m!}} e^{-2x^2} |m\rangle\langle n|, \tag{9.54}$$

$$pr_{\theta_i} = \mathbf{Tr}\left[\hat{\Pi}_i \hat{\rho}\right] \tag{9.55}$$

and $\mathbf{N}$ is a normalisation. The expression for Π_i comes from the overlap $\langle\theta, x|n\rangle$ between the eigenstates for the quadrature and number operators.

Figure 9.8 (a-d) show the density matrices for two different coherent coherent states before and after storage in the memory. Once the density matrix is computed various properties of the states such as the photon number statistics, shown in fig. 9.8 (e,f), can be computed. We find that the even after storage, the photon number statistics are Poissonian. The statistics expected for an optimal classical memory, one that adds one unit of vacuum noise, are shown for reference. The fidelity $F = \mathbf{Tr}\left[\sqrt{\sqrt{\rho_{in}}\rho_{out}\sqrt{\rho_{in}}}\right]^2$ between the input and output states can be calculated to quantify the effect of the memory on stored states. This, however, provides a state-dependent measure of the memory performance. To estimate the effect of the memory on arbitrary input states we need to use a state-independent benchmark, such as the $T - V$ benchmark.

9.4.2 T-V criterion

The memory performance can be characterised by how well the phase and amplitude of an input signal are preserved. These quantities can be degraded through loss, phase instability or the addition of noise photons through non-linear processes. The later of these, is less easily remedied by technical improvements to the memory set-up and is of the most concern. We therefore characterise the signal reduction and the added noise for a variety of input states and show that the added noise is sufficiently small that the memory is operating in a quantum regime.

We use the $T - V$ representation [a76, a77] for benchmarking our memory, which was adapted from use in characterising teleportation experiments in [a78]. This considers a measure of the preservation of signal-to-noise ratio T and added quantum noise V in a system. T and V defined as

$$T = T^+ + T^-, \tag{9.56}$$

$$T^{\pm} = \frac{\mathbf{R}^{\pm}_{out}}{\mathbf{R}^{\pm}_{in}}, \tag{9.57}$$

$$V = \sqrt{V^+_{CV} V^-_{CV}} \tag{9.58}$$

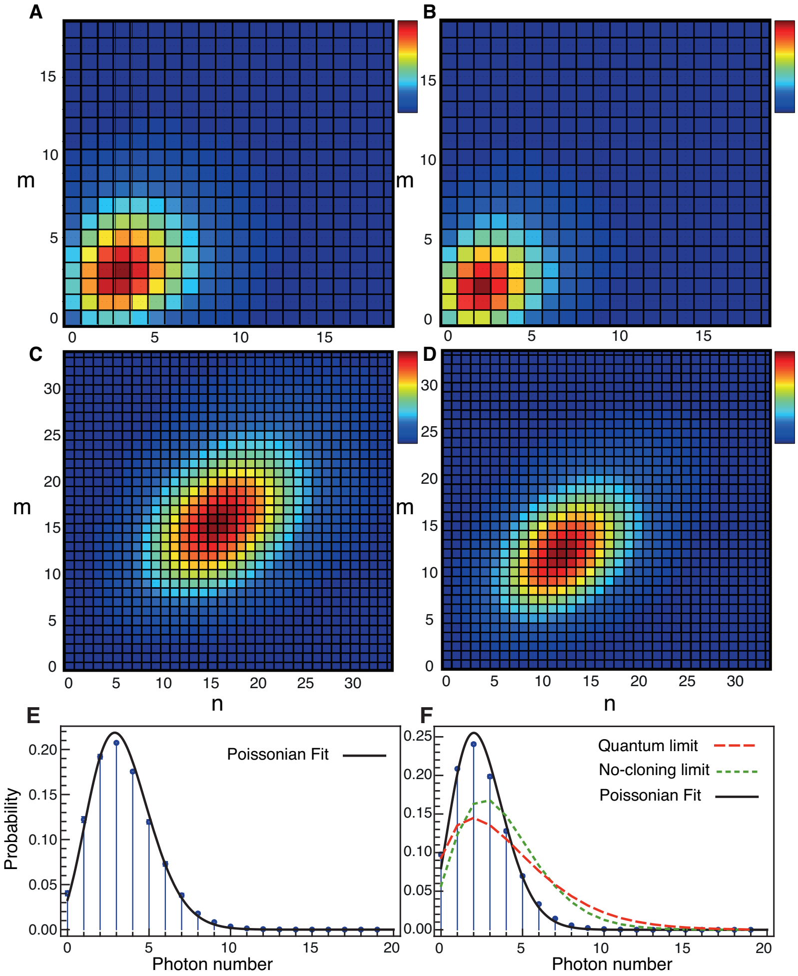

Figure 9.8: Density matrix elements for two sets of input and output coherent states, $n \approx 3$ **a, b** and $n \approx 16$ **c, d** Photon number distribution for input and echo pulses, respectively. The black line shows the fitted Poissonian distribution. The green dotted line represent the no-cloning limit and the dashed red line shows the boundary for the quantum limit.

where **R** is defined by

$$\mathbf{R}^{\pm} = \frac{4(\alpha^{\pm})^2}{V^{\pm}}. \tag{9.59}$$

$V_{CV}^{\pm}$ are conditional variance measurements that would require access to an entangled input state. We can however, estimate V from the noise added to a separately characterised input state using the definition

$$V = 1 - \eta + V_{Noise} \tag{9.60}$$

where η is the memory efficiency and V_{Noise} is the added noise.

Figure 9.7 (d) shows the variances of the input and output for a memory experiment with 74% efficiency. Because the input state has technical phase noise, there is a risk of underestimating the noise added by the memory because it may be masked by the technical noise. Indeed, the output state has less noise than the input. To be conservative, we assume that all of the noise on the output state has been added by the memory, which will underestimate the memory performance.

Figure 9.9 shows the $T - V$ parameters (a) along with the fidelity between the input and output states (b) for a number of different input state amplitudes. The shaded region on the bottom right corresponds to the no-cloning regime where the memory output is the best copy of the input state. These were adjusted to account for detection inefficiency, which could lead to an underestimate of the added noise. Even with this conservative approach, the memory convincingly demonstrates operation in the no-cloning regime.

The data used for these plots were taken with the memory providing an average efficiency of 78%. In fig. 9.9 (b), the solid line is a theoretical curve [a79] for the fidelity assuming an ideal coherent state at the input and a noiseless, 78% efficient memory. The phase noise added in our experiment causes the measured fidelities to drop well below this line for larger coherent states, however, low amplitude coherent states are less sensitive to phase noise.

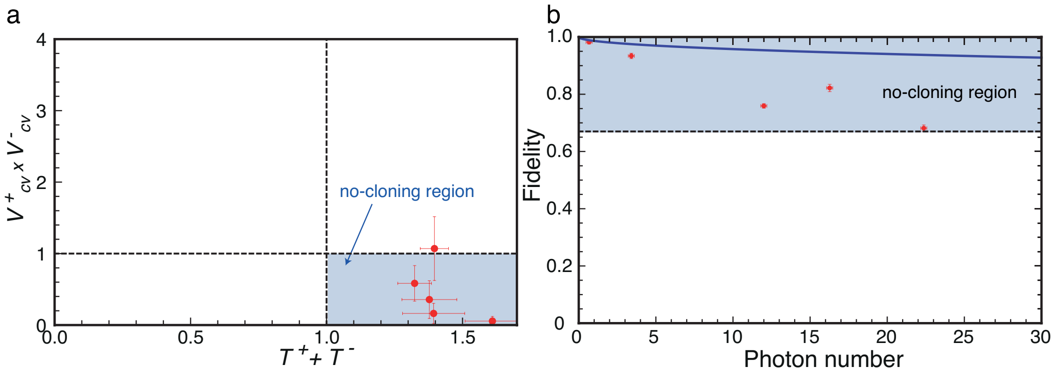

Figure 9.9: a) Diagram representing conditional variances vs transfer coefficient. The shaded rigion on the bottom-left specifies quantum region. The error bars indicate the statistical error of the experimental data. b) Fidelity of the memory for various coherence states with different photon numbers. The solid line is the predicted curve for a 78% efficient memory assuming no noise is added to the coherent state.

CHAPTER 10

Extensions of the GEM protocol

With a demonstration that a three-level gradient echo memory provides high efficiency and high fidelity quantum performance. The rest of the projects related to GEM turned to extending the basic GEM protocol. In addition to the work that I am reporting here, there have been successful efforts by other members of the research team to demonstrate optical pulse sequencing [a11], spectral manipulations of stored pulses [a41] and spatially multi-mode storage [a80].

The first parts of this chapter cover work done on exploring the beam-splitter-like nature of the memory. These include a demonstration of interference between an optical mode and the atomic spin coherence as well as multi-mode interference between frequency-separated optical modes and the atomic coherence. I show that the interference can be extended to perform arbitrary unitary operations by using a series of memories. The final section presents the storage of two independent frequency modes by writing them to different Zeeman levels.

10.1 Interference effects in GEM

As described in chapter 9, the read and write stages of the memory can be treated as being analogous to a beam-splitting operation acting between an optical mode and an atomic coherence. Both the splitting ratio and the interference phase are controlled optically via the strength of the Raman coupling field. While this is a useful model for inefficiencies in the read and write steps [a18], it can be also be used as an actual beam-splitter by interfering an optical mode that is being stored in the memory with a spin coherence that is being recalled to an optical mode.

10.1.1 A time-bin beam-splitter

The conceptual layout for the experiment is illustrated in fig. 10.1. We prepare the atomic coherence by storing a pulse of light in the atomic memory. A second pulse is sent into the memory just as the stored field is being recalled, leading to interference that suppresses or enhances the recall of the stored pulse. As discussed in section 9.2, the effective optical depth d of a three-level GEM is

$$\beta = \pi d \frac{\Gamma}{\eta} \frac{|\Omega|^2}{\Delta^2} \tag{10.1}$$

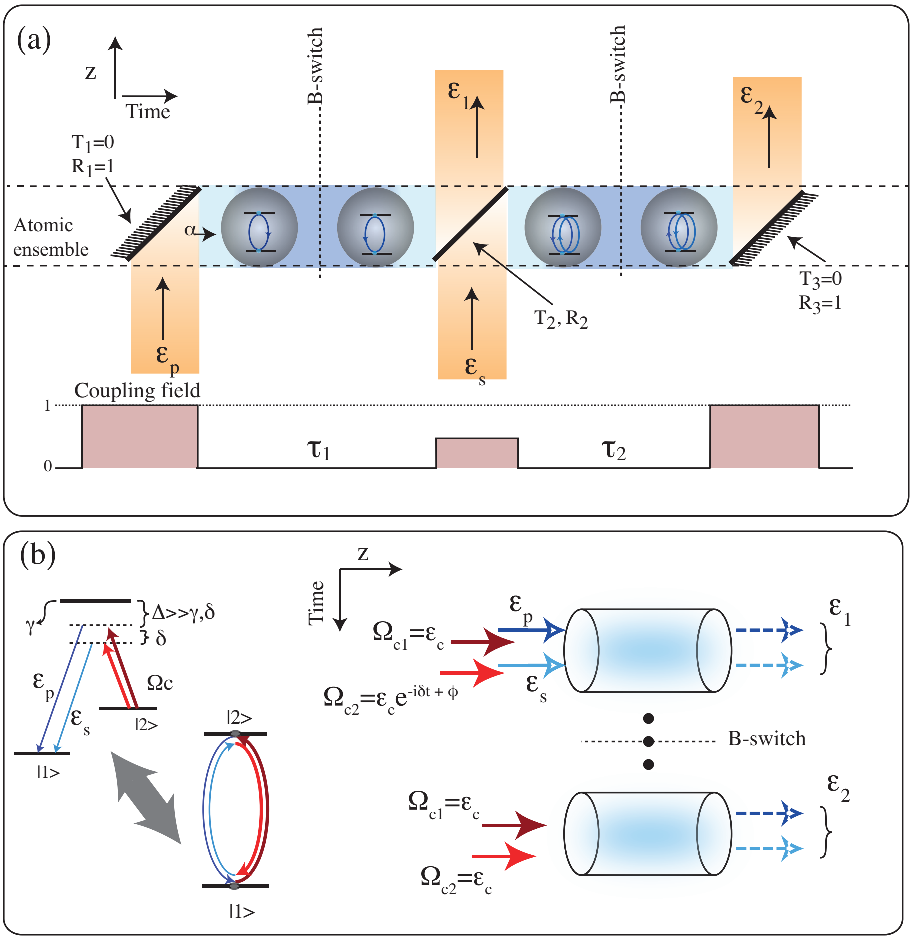

Figure 10.1: (a) Schematic representation of atom-light interference in the memory. The probe pulse, $\hat{\mathcal{E}}_p$, is fully absorbed in the atomic spin coherence (α). The second steering pulse, $\hat{\mathcal{E}}_s$, enters the memory at the precise time that the first echo is being emitted so that it can interfere with the recalled light. The interference is determined by relative phase of the pulses and the effective beam-splitter (T_2, R_2), which is controlled by the strength of the Raman coupling field. The remaining atomic coherence can be recalled later as $\hat{\mathcal{E}}_2$. (b) Left: Double-Λ level structure and optical fields used for interference of two Raman absorption paths of signal fields (probe and steering) with different frequencies. Both Λ transitions drive the same coherence. Right: The procedure used for observing double-Raman interference. The probe and steering pulses are sent into the memory each with a corresponding coupling field. Interference between the unabsorbed probe pulses, $\hat{\mathcal{E}}_1$, and the atomic coherence, which is recalled from the memory as $\hat{\mathcal{E}}_2$, can be observed by varying the relative phase of the two Λ transitions.

For the writing stage the transmissivity, $T(\beta)$, of the effective beam-splitter is the fraction of the input light field that is leaked through the memory so that $T(\beta) = e^{-2\pi\beta}$, while the fraction of the light written into the memory is given by the reflectivity $R(\beta) = 1 - T(\beta)$. For the reading stage, the $R(\beta)$ will be the fraction of the polariton that is converted into a recalled optical field while $T(\beta)$ will be the fraction that remains in the memory.

Because $T(\beta)$ and $R(\beta)$ are defined by the strength of the coupling field, the transmittivity of the beam-splitting can be tuned through the power of the coupling field. A series of reading and writing events, as shown in fig. 10.1(a), can then be described using appropriate reflectivites. The amount of light recalled in the first echo is given by $\hat{\mathcal{E}}_1 = \sqrt{R_1 R_2} e^{-\gamma_0 \tau} \hat{\mathcal{E}}_p + e^{i\theta} \sqrt{T_2} \hat{\mathcal{E}}_s$, where $\hat{\mathcal{E}}_p$ is an initial probe pulse and $\hat{\mathcal{E}}_s$ is a second pulse, which we label the steering pulse, that enters the medium at the time that the probe pulse is recalled. The exponential term arises from the decay of $\hat{\mathcal{E}}_p$, at a rate γ_0, during the storage time τ and the phase θ can be chosen at will. This equation shows that interference can arise between recalled fraction of $\hat{\mathcal{E}}_p$ and $\hat{\mathcal{E}}_s$ and, in particular, if $\sqrt{R_1 R_2} e^{-\gamma_0 \tau} = \sqrt{T_2}$ and $\theta = \pi$ then $\hat{\mathcal{E}}_1$ can be fully suppressed. This simple analysis ignores other details such as the matching of the temporal modes of the pulses. Other factors that limit ideal interference will be discussed later when we analyse the results of our experiments.

The polariton picture, section 9.1.1, can be used to visualise the dynamics of the time-domain beam-splitting operation. As the probe pulse is absorbed and the atomic coherence de-phases, the polariton evolves from to higher spatial frequencies and becomes primarily an atomic spin-wave. The position of the mode in k-space represents how much the atoms in the spin wave have de-phased. A higher atomic frequency gradient means there is a larger de-phasing and therefore faster progression to higher values of k. When the magnetic field gradient is reversed, the polariton propagates back towards $k = 0$ at which point the coherence is re-phased and the polariton is converted into an optical field which exits the memory. The $k = 0$ crossing is analogous to the interface of a beam-splitter. At this point the Λ transition couples the optical mode of the probe field, propagating in the (t, z) plane, and the spin coherence of the atomic ensemble, propagating in the (t, k) plane.

Figure 10.2 shows a numerical simulation of the time-domain interference scheme using the methods described in section 6.5. The simulation shows the evolution of the electric field in real space (the (t, z) plane) and the atomic spin coherence in Fourier-space (the (t, k) plane). In this numerical simulation, the phase of the steering pulse is chosen such that a suppression of the echo from the probe pulse is observed. The intensity of the coupling field is chosen so that the effective splitting ratio is roughly 0.5. Constructive interference into the atomic polarisation occurs and the resultant polariton is recalled in field $\hat{\mathcal{E}}_2$ after the second gradient switch.

Figure 10.3 shows the real components of the electric field and the atomic coherence to illustrate the interference for the same simulation parameters. The phase of the amplitude transmission and reflection coefficients in our beam-splitter analogue depend on both the phase of the coupling field, Ω_c, and the phase of the atomic coherence, $\hat{S}(t, k)$. This provides optical control over the relative phase between the two inputs to the interference.

The experiment was performed using an ensemble of warm ^{87}Rb atoms and a linear switch-able varying magnetic field as described in [a3]. The coupling and probe fields are passed through acousto-optic modulators (AOMs) that allow us to create the required pulse sequences by driving them with appropriate RF signals. To generate the probe and steering pulses, the RF signals were created using separate, but phase locked, arbitrary

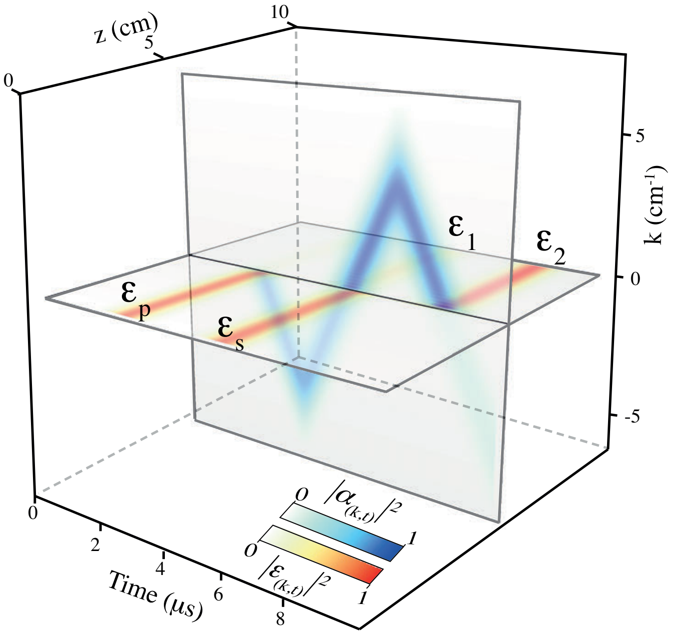

Figure 10.2: Numerical simulation showing interference between the electric field, plotted on the z-t plane, and the atomic coherence, plotted on the k-t plane, where the second light pulse is out of phase from the first echo field. The parameters used in the simulations are: $gNL/\gamma = 40$, $\Omega_c/\Delta = 0.75$, $\Omega_c(t = 4\mu s) = 0.7\Omega_c(t = 2\mu s)$ and $\phi_{\hat{\mathcal{E}}_s} - \phi_{\hat{\mathcal{E}}_p} = \pi$. The magnitudes of the electric field and the atomic coherence are plotted.

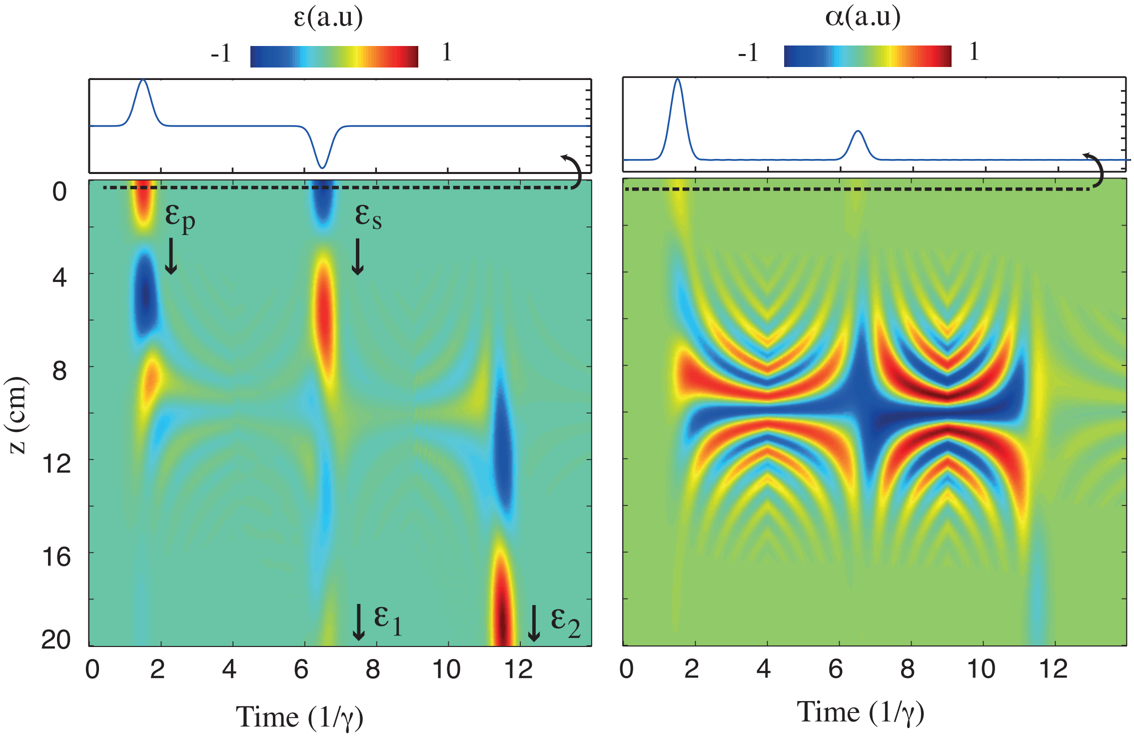

Figure 10.3: The real parts of the electric field (left) and the atomic coherence (right) for a simulation of the time-domain interference experiment. The insets above the plots show the electric field and atomic coherence at the input of the ensemble.

waveform generators and were combined together before the AOM. In this manner, the frequency, phase and amplitude of the coupling, probe and steering fields can be independently controlled. The coupling field power used for maximum coupling between the optical and atomic modes was 330 mW and was adjusted to control the coupling. The probe and steering pulses have the same frequency and were on the order of few μW. The coupling field, blue detuned by 3 GHz from the $S_{1/2}, F = 2 \rightarrow P_{1/2}, F' = 2$ transition, is Raman resonant with the probe and steering pulses which are blue detuned from the $S_{1/2}, F = 1 \rightarrow P_{1/2}, F' = 2$ transition.

We stored a 4 μs probe pulse in the memory and recalled it after a storage time of $\tau_1 = 10\mu s$. The steering pulse was injected just as the atomic coherence excited by the probe returned to $k = 0$. We label the light detected at this time as $\hat{\mathcal{E}}_1$, and integrate the detector signal over the pulse duration to obtain a value for the pulse energy. The polariton that remains in the atomic medium after the first recall is itself recalled after storage time $\tau_2 = 10\mu s$. We detect it in the same manner as $\hat{\mathcal{E}}_1$ and label it $\hat{\mathcal{E}}_2$. Figure 10.1(a) shows the sequence of pulses that are stored, interfered and retrieved along with the coupling field intensity for each step.

The energies of the recalled pulses, $\hat{\mathcal{E}}_1$ and $\hat{\mathcal{E}}_2$, were measured as a function of the relative phase of the probe and steering pulses. The phase of the atomic coherence depends on the relative phase of the coupling and probe fields. It is therefore possible to control the phase of the interference by scanning the phase of either the steering pulse or the corresponding coupling field. Figure 10.4(a) shows interference fringes for $\hat{\mathcal{E}}_1$ (blue, dashed line) and $\hat{\mathcal{E}}_2$ (red, solid line) obtained by varying the phase of the coupling field corresponding to the steering pulse. This was accomplished by varying the phase of the RF signal that drives the coupling field AOM during the interference event relative to its phase during the storage of the probe pulse. For this data, the powers of $\hat{\mathcal{E}}_p$ and $\hat{\mathcal{E}}_s$ were equal and the coupling field power during the interference event was tuned to find the maximum fringe visibility on $\hat{\mathcal{E}}_2$, which was found to be 68%. The visibility of $\hat{\mathcal{E}}_1$ echo, 23%, is substantially lower due to the power mismatch of the steering pulse and the recalled atomic coherence required to optimise the interference in $\hat{\mathcal{E}}_2$. The reflectivity corresponding to the recall of $\hat{\mathcal{E}}_1$ is 37%.

Control over the effective beam-splitter ratio is demonstrated in fig. 10.4 (b). It can be seen that by varying the coupling field power the effective splitting ratio can be tuned to find a maximum in the interference. For this data, the power contained in the steering pulse was adjusted to provide good visibility for both $\hat{\mathcal{E}}_1$ and $\hat{\mathcal{E}}_2$. It is interesting to note that for strong coupling fields, one optical pulse is written into memory while another is being recalled with little interference between the two, analogous to a high-reflectivity beam-splitter. For a weak coupling field, on the other hand, the effective beam-splitter becomes fully transmissive, again meaning no interference between the pulses as the steering pulse passes straight through without storage and the probe pulse remains trapped in the atomic coherence.

10.1.2 A frequency-domain beam-splitter

The second scheme fig. 10.1(b) is a frequency domain interferometer. In this case the atomic coherence is simultaneously driven via two distinct Raman transitions. Interference occurs between the two non-degenerate absorption paths. In this case, the probe and steering pulses are co-propagating and enter the medium simultaneously but are separated in frequency by more than the memory bandwidth. In the far-detuning and adiabatic

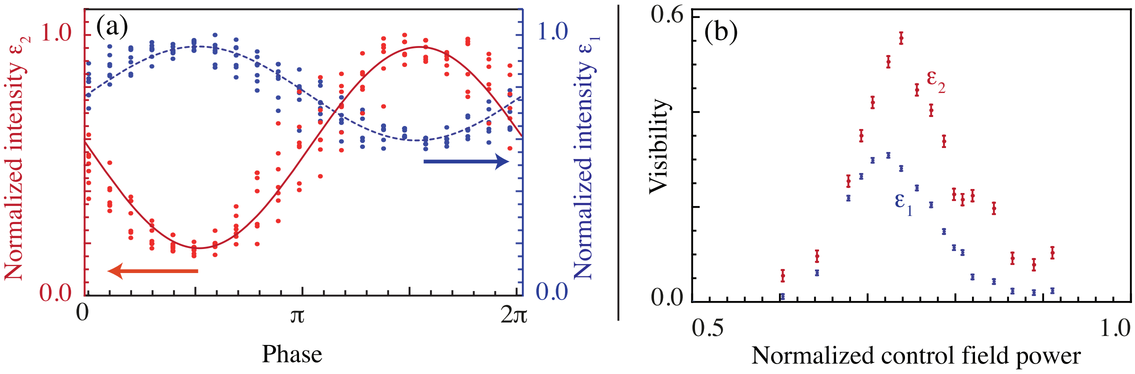

Figure 10.4: (a) Atom-light interference fringes at different times resulted from interaction of the steering pulse with echo generated from the probe pulse. The first arm of the interferometer which is in the optical mode leaves the memory(blue data) and the second arm is stored as an atomic coherence that is transformed back to the light field after re-switching the B-field (red data). The dashed blue and solid red lines are sinusoidal fits to the corresponding data. The red and blue data yield a fringe visibility of 68% and 23%, respectively. (b) Visibility of fringes for two pulses separated in time at the first (blue points) and second (red points) reading stage as a function the normalised coupling field power.

regimes, this double-Λ system is equivalent to a quasi-two-level system interacting with two fields of different Rabi frequency fig. 10.1 (b). The interference between the two Λ transitions will change the response of the medium to the probe and steering pulses. When they interfere destructively, the absorption of the probe and steering fields is suppressed and both pulses are transmitted through the medium. When the two Λ transitions are in-phase, both pulses are coherently absorbed and can be recalled later on-demand.

The properties of the interference can be controlled through the coupling fields. The relative intensity and phase of the two coupling fields control the superposition of the probe and steering pulses that is transferred to the atomic coherence. This effect has been explored in EIT experiments [a81, a82]. Unlike EIT, however, the optical modes that are not coupled to the atomic coherence in the Λ-GEM scheme propagate through the atomic medium with little loss.

The frequency difference between the probe and steering fields was set to 1 MHz, which was larger than the memory bandwidth of 300 kHz to avoid overlap between two broadened Raman lines. Each of the probe and steering fields has a corresponding coupling field which is tuned to Raman resonance. The pulse length, 4 μs, was chosen to give a bandwidth slightly less than the memory bandwidth. Figure 10.5 shows the interference fringe obtained by varying the relative phase between the two Raman absorption lines. This was done by sweeping the phase of one of the coupling fields. The powers of the coupling fields are equal, 160 mW each, and remain constant throughout the storage and retrieval process. $\hat{\mathcal{E}}_1$ is the portion of the probe and steering pulses that does not get stored in the memory and $\hat{\mathcal{E}}_2$ is the portion that is retrieved from the memory after a 10 μs storage period. The energies of $\hat{\mathcal{E}}_1$ and $\hat{\mathcal{E}}_2$ are measured by integrating the detector signal over the pulse period. The lower fringe visibility on $\hat{\mathcal{E}}_1$ can be attributed to inefficient operation of the memory. Light that is not coupled to the atomic spin-coherence does not contribute to the interference fringe and adds an offset to the fringe for the pulse that is not stored.

From an operational standpoint, this second experiment can be thought of as the frequency-domain counterpart to the first. While the first experiment demonstrated a

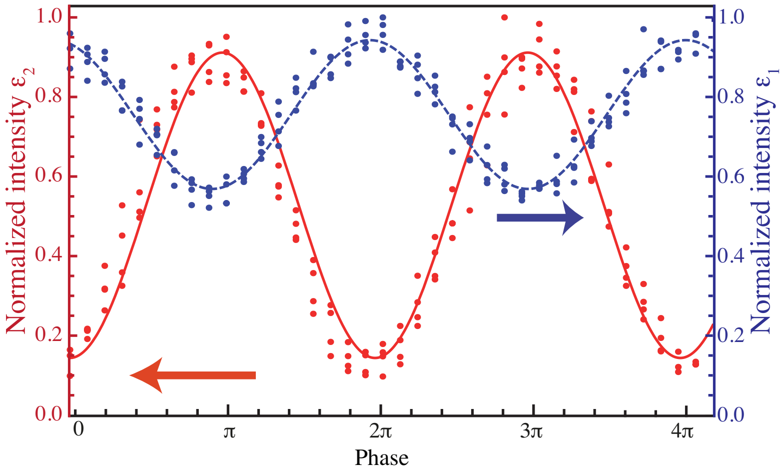

Figure 10.5: Interference fringes from $\hat{\mathcal{E}}_1$ (blue) and $\hat{\mathcal{E}}_2$ (red) that resulted from interference between double-Λ transitions created from the probe and steering pulses of different frequency stored simultaneously in the memory. The dashed blue and solid red lines are sinusoidal fits to the corresponding data. The red and blue data yield a fringe visibility of 73% and 25%, respectively.

beam-splitting operation between two pulses separated in time, the second demonstrates a beam-splitting operation between simultaneous pulses separated in frequency.

In both the time and frequency domain interference experiments, the fringe visibility for the stored coherence, recalled as $\hat{\mathcal{E}}_2$, is less than unity. We attribute this primarily to spatial and temporal mode mismatch between the probe polariton and the steering pulse. We believe that this is mainly due to the atomic motion and non-zero transverse magnetic field, which affects the echo signal for long storage times. This can be justified by the larger visibility measured in the frequency-domain interference scheme, where interference occurs between pulses that simultaneously interact with the atomic coherence. During the storage time, atomic diffusion can change the spatial mode of the coherence and as a result, the echo signal will have a slightly different mode compared to the input signal. This effect is negligible for shorter storage times. The presence of a transverse magnetic field can induce an extra spatial frequency (k_x and k_y) during storage. The transverse k vector is imprinted to the echo signal at the readout diverting the output optical field slightly compare to the steering pulse. In the beam-splitter analogy, this amounts to a poorly aligned interferometer. An inhomogeneous longitudinal magnetic field can alter the shape of the echo signal compared to its input leading to temporal mode mismatch. We anticipate, therefore, that the visibility could be improved by increasing the buffer gas pressure or using a cold atomic sample in order to increase the time of flight of the atoms and taking extra care with the magnetic environment to prevent pulse deflection and distortion. For the time-domain interference, numerical simulations reveal that, in the limit of large optical depth, the interference visibility of the system can approach unity for both interferometer outputs.

10.2 Linear operations in memory

The results of the previous section hint at a broader range of operations that can be performed using the memory. Here, we propose a novel architecture for manipulating optical quantum information by using quantum memories, not just for the synchronisation and buffering of optical states, but also to perform linear operations on them. It is known, for example, that a network of arbitrary beam-splitters is capable of performing a arbitrary linear operations. An operation on N optical modes requires $^{N}C_2$ beam-splitters and phase-shifters [a83], giving a quadratic scaling in N. This forms the basis of schemes to develop photonic networks that are capable of performing arbitrary quantum computations with conditional measurements. The time-domain beam-splitter presented in the previous section provides a method to perform arbitrary operations on optical modes that are separated in time.

Figure 10.6 shows the conceptual layout for performing arbitrary operations on time-domain optical modes. A series of N memories is used to store N input pulses, each operating with a different control field amplitude. By selecting reflectivities of the memories, it is possible to map the input pulses into the memories according to an arbitrary unitary operator.

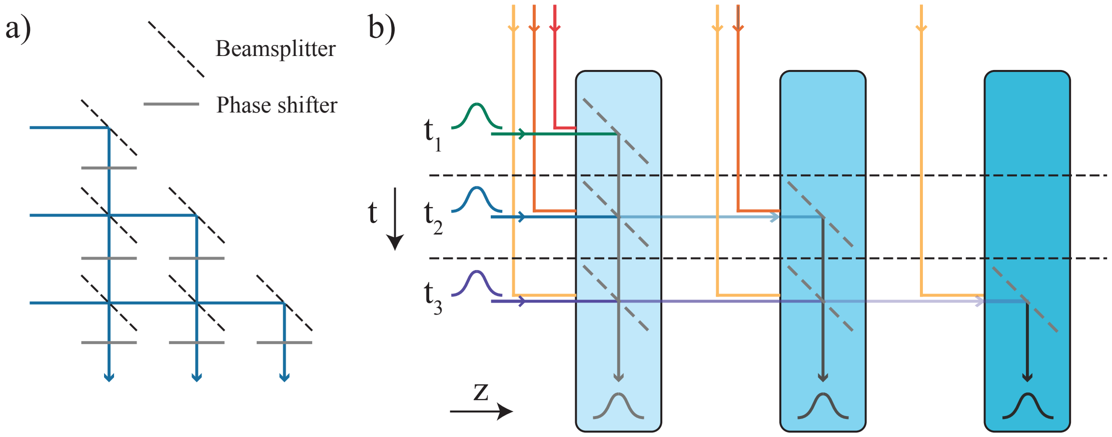

Figure 10.6: a) Arbitrary unitary operations on optical modes can be performed with $^{N}C_2$ beam-splitters ans phase-shifters. b) A scheme for implementing arbitrary operations on time-bin q-bits in memory.

10.2.1 Frequency-domain operations with multiple excited states

The frequency domain provides a similar opportunity for performing operations, however the mechanism is slightly more subtle. In this case, pulses are separated in frequency rather than time so mapping arbitrary superpositions of the input pulses occurs as a result of control fields that are also separated in frequency rather than time. In this protocol, we consider a three-level system that includes multiple signal fields which are separated in frequency. The atomic structure, depicted in fig. 10.7 (a), consists of a ground state, $|g\rangle$, a meta-stable state, $|s\rangle$, with de-phasing rate δ and N excited states, $\{|e_1\rangle, \ldots, |e_N\rangle\}$. Each of the signal fields, $\hat{\mathcal{E}}_k$, is coupled to the spin excitation by a corresponding bright coupling field Ω_k at a detuning Δ_k, resulting in multiple Raman-Λ transitions that drive the spin excitation of the atomic ensemble. Each excited state is assumed to have the same decay rate, Γ, and each of the Raman transitions the same two-photon detuning, δ.

In the far-detuned regime, $\Delta_k \gg \beta\Gamma$ where β is the optical depth [a57], we can adiabatically eliminate the excited states by assuming that $\partial_t \hat{\sigma}_{gk} \approx 0$. Making dipole, rotating-wave and pure-state approximations as described in chapter 6, the Heisenberg equations

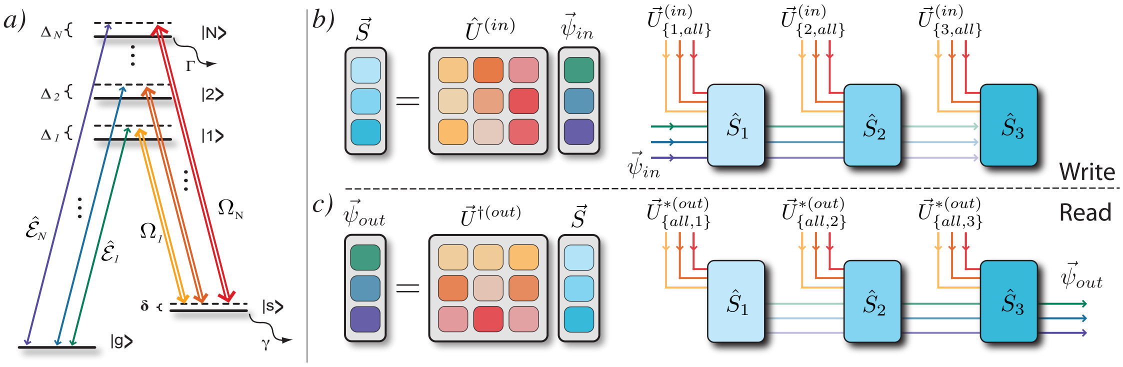

Figure 10.7: a) The multi-Λ level structure used in the protocol consists of a ground-state $|g\rangle$, a meta-stable state $|s\rangle$ and N excited states, $|e_k\rangle$. Optical information carried in the probe modes, $\hat{\mathcal{E}}_k$, can be mapped onto the spin-coherence between $|g\rangle$ and $|s\rangle$ via a two-photon Raman scattering mediated by the coupling fields, Ω_k. b) Controlling the coupling field amplitudes results in the mapping of a unique superposition of the probe modes into each of a set of memories that are placed in series along the optical path. By selecting appropriate coupling field amplitudes $\vec{U}^{(\text{in})}_{\{i,\text{all}\}}$ (represented by column vectors), an arbitrary unitary operation $\hat{U}^{(\text{in})} = [\vec{U}^{(\text{in})}_{\{1,\text{all}\}}, \vec{U}^{(\text{in})}_{\{2,\text{all}\}}, \vec{U}^{\text{in}}_{\{3,\text{all}\}}]^T$ is implemented during the mapping from frequency-separated optical modes to spatially separated spin-waves. c) A second arbitrary unitary operation $\vec{U}^{\dagger(\text{out})} = [\vec{U}^{*(\text{out})}_{\{\text{all},1\}}, \vec{U}^{*(\text{out})}_{\{\text{all},2\}}, \vec{U}^{*(\text{out})}_{\{\text{all},3\}}]$ can applied during the read phase of the scheme.

of motion for the probe field and collective atomic operators [a40] can then be written

$$\partial_\tau \hat{S} = -(\gamma' + i\delta')\hat{S} + i\sqrt{d}\Gamma \sum_{k=1} \frac{\Omega_k^*}{\Delta_k}\hat{\mathcal{E}}_k, \tag{10.2a}$$

$$\partial_\xi \hat{\mathcal{E}}_k = i\sqrt{d}\frac{\Omega_k}{\Delta_k}\hat{S}, \tag{10.2b}$$

where d is the optical depth. The two-photon detuning, $\delta' = \delta - \sum_k(|\Omega_k|^2/\Delta_k)$, and de-phasing rate of $|S\rangle$, $\gamma' = \gamma + \Gamma\sum_k(|\Omega_k|/\Delta_k)^2$, have been adjusted to account for the ac-Stark shift and power broadening, respectively. The dispersion for each field is incorporated into the envelope functions with the transformations $\hat{\mathcal{E}}_k \rightarrow \hat{\mathcal{E}}_k e^{-id\xi\Gamma/\Delta_k}$, and $\Omega_k \rightarrow \Omega_k e^{-id\xi\Gamma/\Delta_k}$ prior to transforming to a moving reference frame. Each signal will acquire a phase of $\beta\Gamma/\Delta_k \ll 1$, resulting in a phase mismatch if the detunings for each Λ system differ. This phase mismatch is small in the far-detuned regime, but could also be eliminated with the introduction of a small angle.

The protocol is based on the observation that the second term on the right side of eq. (10.2a) describes the projection of the input modes onto the vector defined by the amplitudes of Ω_k/Δ_k. We can define a transformed set of modes $\hat{\mathcal{E}}'_j = \sum_{k=1}^N U_{jk}\hat{\mathcal{E}}_k$ that includes the mode $\hat{\mathcal{E}}'_c$, which is in turn defined by

$$\hat{\mathcal{E}}'_c \equiv \frac{1}{\tilde{\Omega}}\sum_{k=1}^{N}\frac{\Omega_k^*}{\Delta_k}\hat{\mathcal{E}}_k \tag{10.3a}$$

$$\tilde{\Omega} = \sqrt{\sum_{k=1}^{N}|\frac{\Omega_k}{\Delta_k}|^2}. \tag{10.3b}$$

In this transformed basis, eqs. (10.2) reduces to

$$\partial_{t'}\hat{\sigma}_{gs} = -(\gamma' + i\delta')\hat{\sigma}_{gs} + i\sqrt{d}\Gamma\tilde{\Omega}\hat{\mathcal{E}}'_c, \tag{10.4a}$$

$$\partial_{z'}\hat{\mathcal{E}}'_c = i\sqrt{d}\tilde{\Omega}\hat{\sigma}_{gs}, \tag{10.4b}$$

$$\partial_{z'}\hat{\mathcal{E}}'_{k\neq c} = 0, \tag{10.4c}$$

which describes a single mode, $\hat{\mathcal{E}}'_c$, interacting with an ensemble of two-level atoms with detuning δ', linewidth γ' and effective coupling strength $g\tilde{\Omega}$. The remaining modes, $\hat{\mathcal{E}}'_{k\neq c}$ propagate unimpeded by the atomic ensemble.

Storage of $\hat{\mathcal{E}}'_c$ results in the mapping $\sum_{k=1}^{N} U_{ck}\hat{\mathcal{E}}_k \rightarrow \hat{S}$, where $\hat{S}$ is a collective spin excitation of $\hat{\sigma}_{gs}$. The application of an arbitrary unitary operation can be accomplished by placing N memories in series along the optical path. The mapping from the input optical state to the memories, illustrated in fig. 10.7 (b), will be $\sum_{k=1}^{N} U_{jk}^{(\mathrm{in})}\hat{\mathcal{E}}_k \rightarrow \hat{S}_j$ if the coupling field amplitudes in the j-th memory are set such that

$$U_{jk}^{(\mathrm{in})} = \frac{1}{\tilde{\Omega}}\frac{\Omega_k^*}{\Delta_k}. \tag{10.5}$$

Similarly, on recall, the mapping $\sum_{j=1}^{N} U_{kj}^{*(\mathrm{out})}\hat{S}_j \rightarrow \hat{\mathcal{E}}_k$ is performed, as illustrated in fig. 10.7 (c), with the coupling fields defined as in eq. (10.5). This allows two independent unitary operations to be performed: one during the write phase, with the resultant state being stored across the set of memories, and one during the read phase, with the result of the operation being restored to a frequency-multiplexed optical state.

The system scales favourably with an increasing number of modes: it requires N memories for a rank-N operation. The protocol does require a total N^2 optical coupling fields; however, at large overall detunings the coupling fields in each memory co-propagate and could be generated by a single electro-optical element. Furthermore, the total power of the coupling fields in each memory is independent of the number of modes. Efficient storage of $\hat{\mathcal{E}}'_c$ in eqs. (10.4) requires a large effective optical depth of the Raman transition, $\beta_{\mathrm{eff}} = \beta\tilde{\Omega}^2\Gamma\gamma^{-1}$, where $\beta = g\mathcal{N}L/\Gamma$ is the resonant optical depth. The effective transition strength depends on the *total* optical power, scaled to the inverse detuning, in all of the coupling fields. An additional advantage is that loss due to memory inefficiency is incurred only once for each mode and, provided that absorption due to the far-detuned excited state is negligible, the resulting overall efficiency is independent of the number of modes.

10.2.2 Analysis with a excited state

The requirement that each Λ transition correspond to a separate excited state is likely prohibitive of a practical implementation of the protocol. It is therefore of interest to determine under what conditions it is possible to use multiple Λ transitions that are detuned from a single excited state. In the limit where the transitions are well separated, each Λ system can be adiabatically eliminated independently and we recover eqs. (10.2). The condition for this to be valid is that $\hat{S}$ must evolve slowly with respect to the frequency separation of the modes so we examine how it is influenced by interactions between the Λ systems.

We proceed by adiabatically eliminating the excited state under the assumption that the detuning is the fastest time-scale of the system: $\beta\Gamma \ll \Delta$ and $|\Delta_k - \Delta| \ll \Delta$, where the detunings are centred around a mean frequency, Δ. In this limit, the equations of motion for a probe field propagating through an ensemble of three-level atoms are:

$$\begin{aligned}\partial_{t'}\hat{S} = &-\left(\gamma + i\delta + i(\Delta - i\Gamma)(|\Omega|^2/\Delta^2)\right)\hat{S}\\ &+ i\sqrt{d}(|\Omega|/\Delta)\hat{\mathcal{E}}_c, \end{aligned} \tag{10.6a}$$

$$\partial_{z'}\hat{\mathcal{E}}_c = i\mathcal{N}(|\Omega|/\Delta)\hat{\sigma}_{gs}, \tag{10.6b}$$

where a $\hat{\mathcal{E}}_c$ is a unitary transformation of the probe field, $\hat{\mathcal{E}}_c \equiv \Omega^*|\Omega|^{-1}\hat{\mathcal{E}}$. The coupling field, $\Omega = \sum_k \Omega_k \exp[i(\Delta_k - \Delta)t']$, and the probe field, $\hat{\mathcal{E}}_c$, contain rapidly oscillating components. We solve eqs. (10.6) and examine the contribution of the rapidly varying terms.

For simplicity, we first solve eq. (10.6a) in the absence of a probe field with the result:

$$\hat{\sigma}_{gs} = \alpha e^{-(\gamma'+i\delta')t'} \prod_{j,k\neq j} \exp\left(-i\frac{\Omega_k^*\Omega_j}{\delta_{k,j}}\frac{\Gamma - i\Delta}{\Delta^2}e^{i\delta_{k,j}t'}\right), \tag{10.7a}$$

$$\approx \alpha e^{-(\gamma'+i\delta')t'}\left(1 - \sum_{j,k\neq j}\frac{\Omega_k^*\Omega_j}{\Delta\delta_{k,j}}e^{i\delta_{k,j}t}\right), \tag{10.7b}$$

where α is an integration constant, $\delta_{k,j} \equiv (\Delta_k - \Delta_j)$, and we have assumed that $\Delta \gg \Gamma$. To move from eq. (10.7a) to eq. (10.7b) we have assumed that $|\delta_{k,j}| \gg |\frac{\Omega_k^*\Omega_j}{\Delta}|$ and eliminated higher-order terms. The magnitude of the perturbation will depend on the particular amplitudes of the fields. We can place a conservative requirement on the magnitude of $\delta_{j,k}$ by assuming equal spacings between modes and roughly equal distribution of the coupling field power across the mode spectrum. With these conditions, we require

$$|\delta_{j,k}| \gg \frac{\Delta|\tilde{\Omega}|^2}{\sqrt{N}}. \tag{10.8}$$

The $\sqrt{N}$ is a result of assuming that the amplitudes of the terms in sum do not interfere constructively and can be added in quadrature.

With the oscillating component of $|\Omega(t')|^2$ omitted, we now consider a general solution of eq. (10.6a) that includes the probe field driving terms:

$$\begin{aligned}\hat{S} =& \alpha e^{-(\gamma'+i\delta')t'} + \frac{i\sqrt{d}\Gamma}{\gamma' + i\delta'}\sum_k \frac{\Omega_k^*}{\Delta}\hat{\mathcal{E}}_k\\ &+ i\sqrt{d}\sum_{k,j\neq k}\frac{\Omega_j^*}{\Delta}\hat{\mathcal{E}}_k\frac{e^{i\delta_{k,j}t'}}{\gamma' + i(\delta' + \delta_{k,j})}, \end{aligned} \tag{10.9a}$$

where we have assumed a form of $\hat{\mathcal{E}} = \sum\limits_k \hat{\mathcal{E}}_k e^{i(\Delta_k - \Delta)t'}$ for the probe field. Again, the magnitude of the rapidly oscillating term in $\hat{S}$ will depend on the amplitudes of $\{\Omega_k^*\hat{\mathcal{E}}_k\}$. A conservative restriction would be

$$|\delta_{k,j}| \gg \sqrt{N} \times \{\gamma', |\delta'|\}. \tag{10.10}$$

This requires that the frequency separation of the modes be large compared to the memory bandwidth.

An implementation of the protocol that uses cold-atom GEM is advantageous because the coupling fields in each memory do not need to co-propagate with the signal fields but can be offset by a small angle. This allows a single memory to be subdivided into segments, each of which is addressed by its own set of coupling fields, without the need for beam-splitters between the segments.

To place estimates on what could be achieved with our present experiment, we use optimistic but plausible parameters. We assume a maximum available optical depth of 1000 in a 5 mm long ensemble. This is somewhat higher than has so far been observed for a single Zeeman sub-level, but could be achieved with the current magneto-optical trap with appropriate optical pumping. We then consider an implementation using a 10 μs pulse duration, a detuning of 250 MHz and a coupling field power of 350 mW. The pulse duration and detuning are the same as for the existing experiment and there is ample pump power available to deliver 350 mW to the ensemble.

Dividing the memory into segments, each with an optical depth of 100, would allow a ten mode implementation of the protocol with modes separated by 50 MHz which satisfies eq. (10.10), the more stringent of the two inequalities, by a factor of 167. This would require a total coupling field power of 32 mW and, in principle, would be $\approx 95\%$ efficient. The total bandwidth, 500 MHz, is less than the bandwidth of available electro-optic modulators, meaning that all of the coupling fields could be generated with only ten electro-optic elements. Dispersion between signal fields in this example would cause a $\approx \lambda/4$ difference between the nearest and furthest detuned modes, requiring a 1.6° spread in the coupling field angles to correct. Increasing the base detuning to 750 MHz, however, would reduce the phase error due to dispersion to $\approx \lambda/20$ without any angle between the coupling fields in each segment.

10.2.3 Numerical simulations of practical gates

Numerical simulations [a59, a84], as described in section 6.5, were performed for similar parameters to verify the performance of the protocol, although a mode spacing of 15 MHz rather than 50 MHz was used to reduce computation time. A further optimisation of the simulation code was to separate out the complex exponential functions that define the carrier frequencies for the control and probe fields. These were pre-computed and stored in vectors to avoid re-calculating the computationally expensive functions for each step of the numerical integration. The results demonstrated an overall efficiency of $(91.2 \pm 0.2)\%$ with an overlap of (0.988 ± 0.003) between the simulated and ideal outputs after two operations.

A ten-mode implementation of the protocol would be sufficient to perform non-trivial linear gates. As a demonstration of the versatility, we simulated an operation that is useful to linear optical quantum computing; the implementation of a conditional cz gate as described in [a22]. Figure 10.8 illustrates how this gate can be performed on a pair of dual-rail qubits using conditional measurements and feeding-forward of measurement results. The left side of the figure (a) shows the network of one- and two-mode linear operations as originally proposed. The right side (b) shows the same unitary operations as implemented in our protocol using a set of eight memories.

I refer the reader to [a22, a23] for a description of the linear network shown in fig. 10.8 (a) and focus only on how the network can be translated into the memory-based protocol. In our proposed implementation, four memories are conditionally prepared by the operation $\hat{U}_1$, acting during storage, and a detection of the state of the other four modes, performed by a read-out step, $\hat{U}_2$. If the operation has been successful, the state of interest

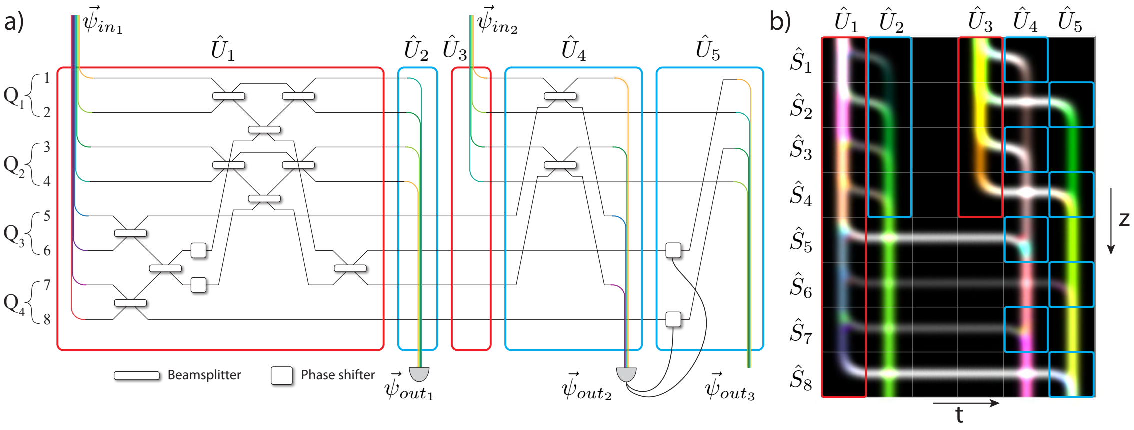

Figure 10.8: A numerical simulation of the protocol implementing the unitary operation required to perform a conditional c-z gate as described in [a22]. (a) The network of linear optical elements as originally proposed. The qubits are path-encoded photons with $\hat{Q}_n$ representing the nth qubit. In this diagram, photons are injected from the top ($\hat{\psi}_{in_n}$) and read out from the bottom ($\hat{\psi}_{out_n}$) for consistency with the simulation. (b) The simulation results, illustrating the propagation of fields through the memory elements as each unitary operation is performed. Each optical frequency is represented by a different color and atomic spin waves are represented as white. The optical modes co-propagate while the spin waves remain spatially separated, each horizontal segment, labelled $\hat{S}_1$ through $\hat{S}_N$, representing a memory element. Each vertical segment is a temporal window in which a storage or recall event can occur to map information between the (frequency-encoded) optical basis and the (path-encoded) spin-wave basis. Every area in the grid therefore has an independently controllable set of coupling field amplitudes. The optical fields enter from the top and propagate rapidly along the spatial dimension until they are stored in a memory cell. On recall, they continue to propagate and exit at the bottom of the grid. The areas outlined in red are storage operations and the areas outlined in blue are recall operations.

is then written into the memories and the successful gate result teleported onto modes of interest via another linear operation, $\hat{U}_4$, and a conditional phase-shift of the remaining qubits, $\hat{U}_5$.

This example highlights some of the advantages of this protocol. The combination of an optically configurable operation with integrated memory provides a simple path to measurement-based non-linearities. Operations involving a large number of modes can be performed in a single step with the result being stored in memory for use at a later point in a computation. The protocol is spatially single-mode, providing inherent interferometric stability and only one type of element is required, eliminating interfaces between devices. The major drawback, replacing simple optical elements with atomic ensemble-based memories, is mitigated by the fact that networks of linear optics would likely require similar memories for synchronisation and to implement the feeding-forward of measurement results.

10.3 Dual-GEM

Of particular relevance for quantum memory is the ability to store qubits with a high fidelity. To this end, we worked on a memory for dual-rail qubits that are separated in frequency. In this approach, a qbit would be encoded as a photon in a superposition of two central frequencies as opposed to a superpostion of a photon or vacuum. This approach is a variation of path-encoding for optical quantum information in which a qbit is encoded based on its presence in one of two pathways. These encodings have the advantage of being insensitive to loss; provided that a photon is detected, loss will not affect the value of the qbit.

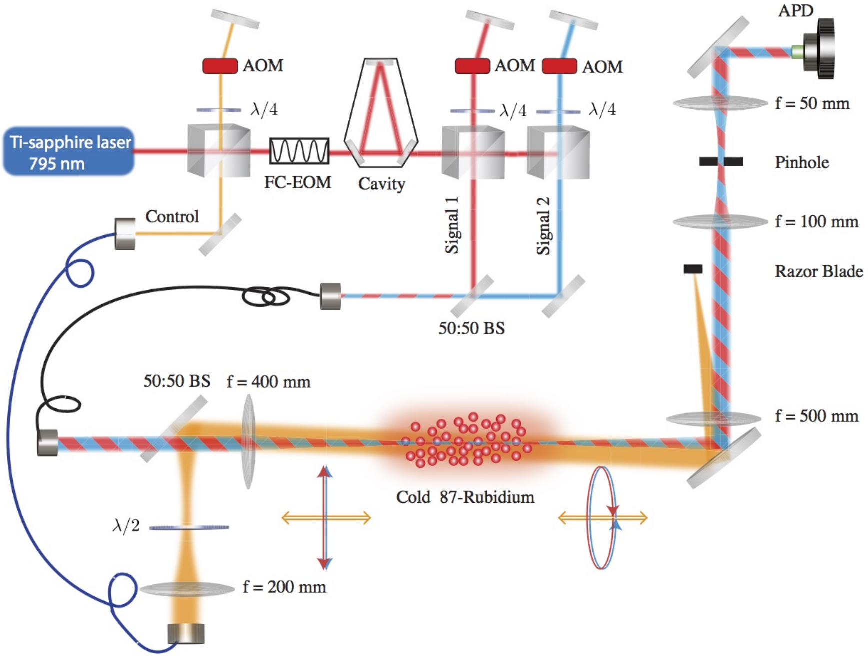

Figure 10.9: The experimental layout for dual-GEM. The two signal fields, each generated from a separate AOM, are prepared in the same linear polarisation before the memory. The control field is in the orthogonal linear polarisation. At the output, the signal field polarisations have been slightly altered, reducing the visibility between them.

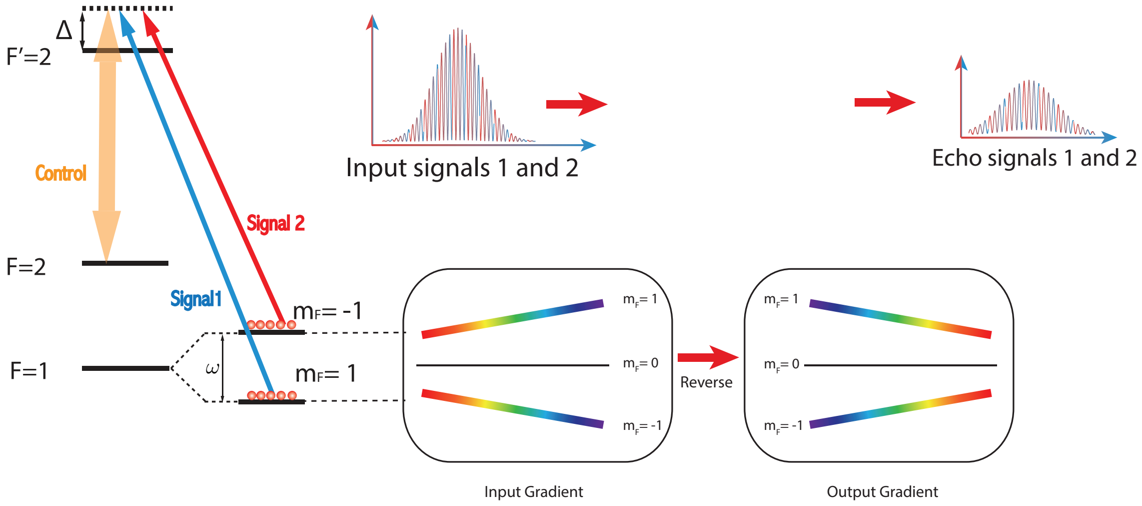

Figure 10.10: A simplified energy diagram of atomic levels illustrating the dual-channel memory.

We demonstrated that GEM can be used as a memory for such frequency encoded states by storing and recalling two-coloured light into two independent spin-coherences in the Zeeman sub-levels of the Rb atom. The atomic frequencies are manipulated by applying tailored magnetic field along the optical propagation axis. A bias magnetic field B_0 lifts the degeneracy of the Zeeman sub-levels by introducing a Zeeman shift given by $\Delta E = m_F g_F \mu_B B_0/h$, where m_F is the Zeeman level number, g_F is the Landé factor for hyperfine state F, μ_B is the Bohr magneton and h is the Planck constant. For ^{87}Rb atoms, the Landé factors are respectively $g_1 = -1/2$ and $g_2 = 1/2$. A single coupling field detuned from the $F = 2 \rightarrow F' = 2$ transition will then produce three Raman absorption lines for a weak signal field on the $F = 1 \rightarrow F' = 2$ transition, corresponding to signal field detunings of 0 (insensitive line) and $\pm\delta_0$ relative to the un-split Raman transition, where δ_0 is given by

$$\delta_0 = \frac{\mu_B}{h} B_0 \simeq 1.4\text{MHz/G} \times B_0. \tag{10.11}$$

To store light using the GEM protocol, a reversible gradient of magnetic field B_z is applied along the ensemble in addition the uniform bias field B_0 to broaden the 2 sensitive lines around $\pm\delta_0$, as illustrated in fig. 10.10.

The level structure of ^{87}Rb results in a complex set of couplings for each of the signal fields; fig. 10.11 shows the Λ systems that are resonant when the signals have a detuning of $\pm 2\delta_B$. This structure could be simplified by using only one circular polarisation for the coupling field and signal fields, however, unequal dipole moments for the transitions associated with each Λ system would result in an imbalance in the storage efficiencies for the two signal fields. We therefore select a linear polarisation for the coupling field so that the two signals are coupled to the memory symmetrically. In this case, the polarisation of each of the signal fields that is coupled to the memory must be calculated.

To analyse the system, we assume that the population rests entirely in the $F = 1$ ground state and that there are no coherences between the Zeeman sub-levels within that manifold. These assumptions are valid provided that the probe field is weak and that the splitting, δ_0, is large compared to the Raman linewidth. We proceed by treating each of the $F = 1$ Zeeman sub-levels independently, starting with the $m_F = 1$ sub-level. Referring to fig. 10.11, signal 1 will produce a coherence, $\hat{\sigma}$, between the $|F = 1, m_F = 1\rangle$ and

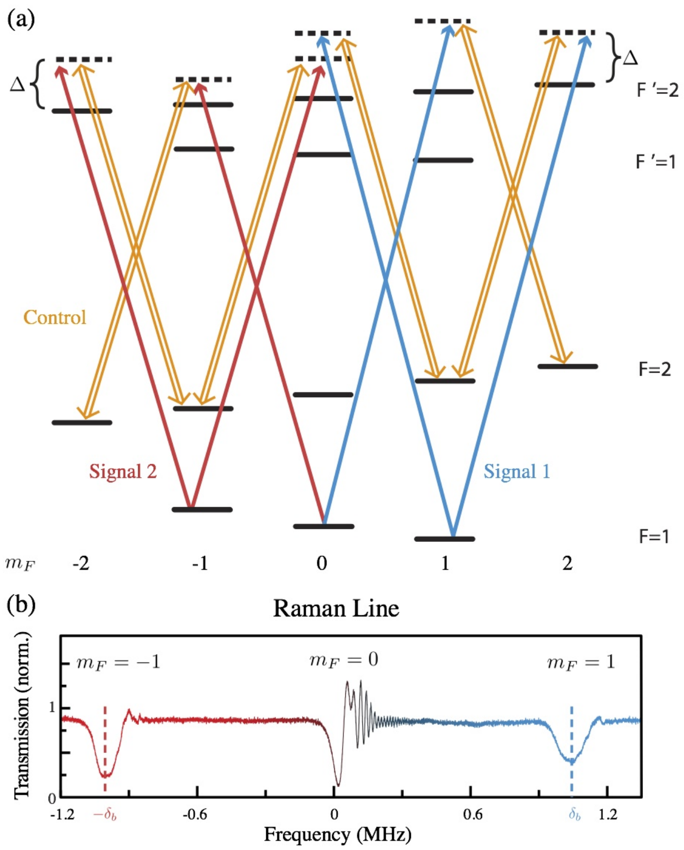

Figure 10.11: a) The energy diagram of atom levels coupled by different laser fields in the scheme. b) Split Raman absorption lines obtained by applying a magnetic field to the atomic ensemble. The red trace shows the lines with a uniform magnetic field and the red trace shows the $\pm 2\delta_B$ lines broadened by an additional gradient the the field. Free-induction decay can be seen in the centre line as well as the un-broadened split lines. This decay is suppressed by the inhomogeneous broadening induced by the gradient magnetic field.

$|F = 2, m_F = 1\rangle$ states via three Λ systems, formed by interaction with the excited states $|F' = 2, m_F = 0\rangle$, $|F' = 1, m_F = 0\rangle$ and $|F' = 2, m_F = 2\rangle$. The equations of motion for signal 1 are

$$\partial_t \sigma = -(\gamma_0' + i\delta')\sigma + ig_{eff}^{-}\hat{\mathcal{E}}^{-} + ig_{eff}^{+}\hat{\mathcal{E}}^{+} \tag{10.12a}$$

$$(\partial_t + c\partial_z)\hat{\mathcal{E}}^{\pm} = ig_{eff}^{\pm} N \hat{\mathcal{E}}^{\pm}, \tag{10.12b}$$

with effective coupling strengths that have the proportionalities

$$g_{eff}^{+} \propto \mu_{(1,1;2,2)}\mu_{(2,1;2,2)} \tag{10.13}$$

$$g_{eff}^{-} \propto \mu_{(1,1;2,0)}\mu_{(2,1;2,0)} + \mu_{(1,1;1,0)}\mu_{(2,1;1,0)}, \tag{10.14}$$

where $\mu_{(m,n;p,q)}$ is the dipole matrix element corresponding to the $|F = m, m_F = n\rangle \rightarrow |F' = p, m_F = q\rangle$ transition. We can further simplify this system by introducing the composite mode

$$g_{eff}\hat{\hat{\mathcal{E}}}_c \equiv (g^-_{eff}\hat{\hat{\mathcal{E}}}^- + g^+_{eff}\hat{\hat{\mathcal{E}}}^+) \tag{10.15}$$

$$g_{eff} \equiv \sqrt{|g^-_{eff}|^2 + |g^+_{eff}|^2}. \tag{10.16}$$

so that equations eqs. (10.12a) and (10.12b) become

$$\partial_t\sigma = -(\gamma_0' + i\delta')\sigma + ig_{eff}\hat{\hat{\mathcal{E}}}_c \tag{10.17}$$

$$(\partial_t + c\partial_z)\hat{\hat{\mathcal{E}}}_c = ig_{eff}N\hat{\hat{\mathcal{E}}}_c. \tag{10.18}$$

The Λ system formed by $|F = 1, m_F = 0\rangle$, $|F' = 2, m_F = 1\rangle$ and $|F = 2, m_F = 2\rangle$ has a small contribution to the memory process due to a weak effective coupling strength; however, because the optical depth associated with this process is ≈ 19 times smaller than that corresponding to the $|F = 1, m_F = 1\rangle \rightarrow |F = 2, m_F = 1\rangle$ coherence, we neglect it in this analysis. The ideal polarisation for storage can then be calculated from the relative amplitudes of g^-_{eff} and g^+_{eff}. We find that it is $0.52|L\rangle - 0.86|R\rangle$, which results in an ideal storage efficiency of 94% for horizontally polarised light.

The results of the dual-channel storage and recall experiment are shown in fig. 10.12. Panels (a) and (b) show the storage and recall of each of the signal fields independently, in the absence of the other. Panel (c) shows storage of the two signal fields simultaneously. The input signal pulse is a Gaussian profile with a 10 μs width and we record reference traces of the inputs (blue) in the absence of the atomic ensemble by blocking the trapping beams and allowing the MOT to dissipate. Traces showing storage and recall (red) are recorded in the presence of the MOT with a flip of the magnetic field gradient occurring immediately after the input pulse is fully stored in the cell. The polarisation of the control field is slightly tuned to balance the efficiencies of the echoes to account for slight polarisation errors caused by the glass vacuum cell and misalignment of the magnetic gradient relative to the optical axis.

The recall efficiencies are 39% and 32% for signal fields one and two, respectively. We demonstrate the dual-rail storage by sending both signal pulses into the memory simultaneously with matched temporal profiles and amplitudes. The resultant combined efficiency is 35% and the interference fringe visibility between the recalled pulses is 73%, compared to an expected visibility of 77%. The relative phase between the echo signals is measured and compared to that of the input pulses and we find that the standard deviation of the echo phase relative to the input is 15°. The distribution of relative phases is shown in fig. 10.12 (d). The distribution of red points shows the 15° standard deviation.

We determined that primary source of the deviation in phase shift was a result of background magnetic field fluctuations. To reduce these, the experiment was triggered to the mains power supply so that the magnetic fields due to mains current would be equal for each run of the experiment. The result was reduction in the phase drift to 6°, as shown by the blue distribution in fig. 10.12 (d); however, at the time those results were taken, our rubidium dispensers were nearing exhaustion, reducing the efficiency and visibility of the output pulses.

The results demonstrate that the gradient echo memory presented can be used to store two signal channels simultaneously and recall them while preserving the relative phase and amplitude of the optical states. Upgrades to our experimental apparatus, in particular enhanced isolation from ambient magnetic fields and a higher optical depth, are

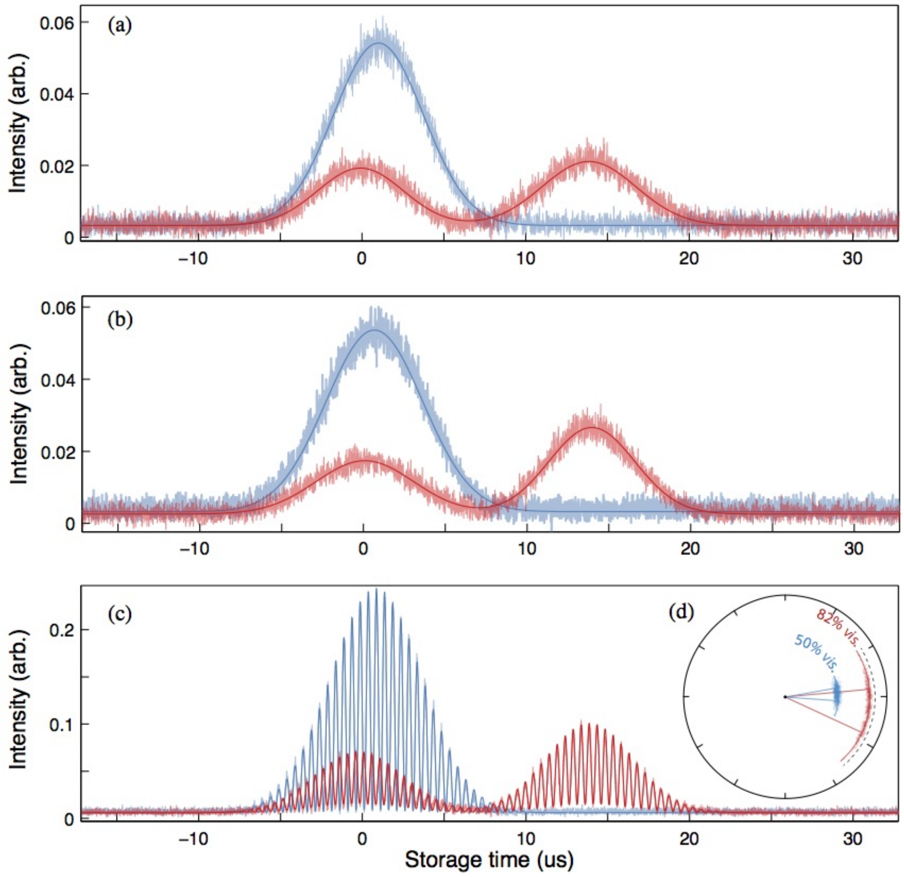

Figure 10.12: Storage and recall of the two signal fields independently (a,b) and both simultaneously (c). The blue traces are the inputs to the memory in the absence of an atomic ensemble and the red traces show the echo pulses as well as light that is leaked through the memory during the storage process.

under-way and promise to improve both the phase stability as well as recall efficiencies.

CHAPTER 11

Conclusions and outlook

The gradient echo memory operating with atomic vapours is the highest efficiency optical quantum memory that has been demonstrated. The memory dynamics are well-described by both analytic and numerical results and, for the cold-atom memory in particular, the experimental results are in excellent agreement with the predicted memory performance.

To demonstrate that the memory is not only efficient but also noiseless, we performed state tomography on the output of the memory. The results confirmed that the gradient echo memory performs as a true quantum memory and that the added noise comes only from technical sources. The tomographic characterisation was done only on the hot-atom memory because of the slow repetition rate for the cold-atom memory. The transient nature of the compression phase of the cold atom trapping results in wait times for the trap to fill and for magnetic fields to decay. The repetition rate is therefore restricted to about two Hertz; too slow to collect the statistical data required for state tomography.

Evidence that the cold atom memory is not susceptible to four-wave-mixing (4WM) noise must therefore come from an understanding of the noise process itself. The electromagnetically induced transparency (EIT) experiments demonstrate that a simple four-level model, one that has been used to theoretically predict the addition of 4WM noise [a65], was in good agreement with the experiment over a wide range of optical depths and in two Rb isotopes. When applied to a gradient echo memory operating in the same system, the model predicts that the contribution from 4WM is negligible.

In addition to characterising the performance of gradient echo memories as single-mode storage devices, we explored extensions to the protocol. We demonstrated interference effects between propagating optical fields and the atomic spin coherence with fringe visibilities of 68% and 73% for time-domain and frequency-domain interference schemes, respectively. These gates could operate on either time-bin or frequency multiplexed qubits or even a combination thereof. The time-delayed beamsplitter scheme could also be used for optical quantum state engineering [a85, a86] and also for optimal Gaussian purification of coherent states from several imperfect copies [a87].

The frequency-domain interference scheme was also demonstrated to integrate the to perform arbitrary unitary transformations into a quantum memory. The operation on frequency multiplexed states makes the protocol naturally suited to work with optical parametric oscillator sources that produce a frequency comb of two-mode squeezed states [a88, a89]. With only a modest increase in optical depth, it should be possible to perform an eight-mode demonstration of the concept. The recent development of a Raman scattering based optical memory in a solid-state ensemble [a90] is also an encouraging step towards compact memories that are suitable for implementation of the protocol.

Bibliography - Quantum Memories

Articles

a1. Lvovsky, A. I., Sanders, B. C. & Tittel, W. Optical quantum memory. *N. Photon.* **3,** 706–714 (2009).

a2. Bussières, F. *et al.* Prospective applications of optical quantum memories. *Journal of Modern Optics* **60,** 1519–1537 (2013).

a3. Hosseini, M., Sparkes, B. M., Campbell, G., Lam, P. K. & Buchler, B. C. High efficiency coherent optical memory with warm rubidium vapour. *Nat. Commun.* **2** (Feb. 2011).

a4. Hedges, M. P., Longdell, J. J., Li, Y. & Sellars, M. Efficient quantum memory for light. *Nature* **465,** 1052–1056 (Apr. 2010).

a5. Sparkes, B. M. *et al.* Gradient echo memory in an ultra-high optical depth cold atomic ensemble. *New Journal of Physics* **15,** 085027 (2013).

a6. Longdell, J. J., Fraval, E., Sellars, M. J. & Manson, N. B. Stopped Light with Storage Times Greater than One Second Using Electromagnetically Induced Transparency in a Solid. *Physical Review Letters* **95,** 2–5. ISSN: 0031-9007 (Aug. 2005).

a7. Dudin, Y. O., Li, L. & Kuzmich, A. Light storage on the time scale of a minute. *Physical Review A* **87,** 031801. ISSN: 1050-2947 (Mar. 2013).

a8. Heinze, G., Hubrich, C. & Halfmann, T. Stopped Light and Image Storage by Electromagnetically Induced Transparency up to the Regime of One Minute. *Physical Review Letters* **111,** 033601. ISSN: 0031-9007 (July 2013).

a9. Saglamyurek, E. *et al.* Broadband waveguide quantum memory for entangled photons. *Nature* **469,** 512–515 (Jan. 2011).

a10. Saglamyurek, E. *et al.* Quantum storage of entangled telecom-wavelength photons in an erbium-doped optical fibre. *Nature Photonics* (2015).

a11. Hosseini, M. *et al.* Coherent optical pulse sequencer for quantum applications. *Nature* **461,** 241 (Sept. 2009).

a12. Fleischhauer, M., Imamoglu, A. & Marangos, J. P. Electromagnetically induced transparency: Optics in coherent media. *Rev. Mod. Phys.* **77,** 633–673 (July 2005).

a13. Marangos, J. P. Electromagnetically induced transparency. *Journal of Modern Optics* **45,** 471–503 (1998).

a14. Fleischhauer, M. & Lukin, M. D. Quantum memory for photons: Dark-state polaritons. *Phys. Rev. A* **65,** 022314 (2 Jan. 2002).

a15. Kraus, B. *et al.* Quantum memory for nonstationary light fields based on controlled reversible inhomogeneous broadening. *Physical Review A* **73,** 020302 (2006).

a16. Gorshkov, A. V., André, A., Lukin, M. D. & Sørensen, A. S. Photon storage in Λ-type optically dense atomic media. III. Effects of inhomogeneous broadening. *Physical Review A* **76,** 033806 (2007).

a17. Moiseev, E. S. & Moiseev, S. A. Scalable time reversal of Raman echo quantum memory and quantum waveform conversion of light pulse. *New Journal of Physics* **15,** 105005 (2013).

a18. Longdell, J. J., Hetet, G., Lam, P. K. & Sellars, M. J. Analytic treatment of controlled reversible inhomogeneous broadening quantum memories for light using two-level atoms. *Phys. Rev. A* **78,** 032337 (2008).

a19. Gorshkov, A. V., André, A., Lukin, M. D. & Sørensen, A. S. Photon storage in Λ-type optically dense atomic media. I. Cavity model. *Phys. Rev. A* **76,** 033804 (3 Sept. 2007).

a20. Duan, L.-M., Lukin, M., Cirac, J. I. & Zoller, P. Long-distance quantum communication with atomic ensembles and linear optics. *Nature* **414,** 413–418 (2001).

a21. Sangouard, N., Simon, C., De Riedmatten, H. & Gisin, N. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics* **83,** 33 (2011).

a22. Knill, E., Laflamme, R. & Milburn, G. J. A scheme for efficient quantum computation with linear optics. *Nature* **409,** 46–52 (2001).

a23. Kok, P. *et al.* Linear optical quantum computing with photonic qubits. *Rev. Mod. Phys.* **79,** 135–174 (1 Jan. 2007).

a24. Nunn, J. *et al.* Enhancing Multiphoton Rates with Quantum Memories. *Phys. Rev. Lett.* **110,** 133601 (13 Mar. 2013).

a25. Boller, K. J., Imamolu, A. & Harris, S. Observation of electromagnetically induced transparency. *Physical Review Letters* **66,** 2593–2596 (May 1991).

a26. Budker, D., Kimball, D., Rochester, S. & Yashchuk, V. Nonlinear Magneto-optics and Reduced Group Velocity of Light in Atomic Vapor with Slow Ground State Relaxation. *Physical Review Letters* **83,** 1767–1770 (Aug. 1999).

a27. Hau, L. V., Harris, S. E., Dutton, Z. & Boulevard, E. H. L. Light speed reduction to 17 metres per second in an ultracold atomic gas. **397,** 594–598 (1999).

a28. Harris, S. E., Field, J. E. & Imamoglu, A. Nonlinear optical processes using electromagnetically induced transparency. *Physical Review Letters* **64,** 1107–1110 (Mar. 1990).

a29. Jain, M., Xia, H., Yin, G. Y., Merriam, A. J. & Harris, S. E. Efficient nonlinear frequency conversion with maximal atomic coherence. *Physical Review Letters* **77,** 4326–4329 (1996).

a30. Wang, Z.-B., Marzlin, K.-P. & Sanders, B. C. Large cross-phase modulation between slow copropagating weak pulses in Rb 87. *Physical review letters* **97,** 063901 (2006).

a31. Appel, J., Figueroa, E., Korystov, D., Lobino, M. & Lvovsky, A. I. Quantum Memory for Squeezed Light. *Phys. Rev. Lett.* **100,** 093602 (9 Mar. 2008).

a32. Honda, K. *et al.* Storage and retrieval of a squeezed vacuum. *Physical Review Letters* **100,** 093601 (Jan. 2008).

a33. Choi, K. S., Deng, H., Laurat, J. & Kimble, H. J. Mapping photonic entanglement into and out of a quantum memory. *Nature* **452,** 67–71 (2008).

a34. Fleischhauer, M. & Lukin, M. D. Dark-State Polaritons in Electromagnetically Induced Transparency. *Phys. Rev. Lett.* **84,** 5094–5097 (May 2000).

a35. Chen, Y.-H. *et al.* Coherent Optical Memory with High Storage Efficiency and Large Fractional Delay. *Physical Review Letters* **110,** 083601. ISSN: 0031-9007 (Feb. 2013).

a36. Boyd, R. W., Gauthier, D. J., Gaeta, A. L. & Willner, A. E. Maximum time delay achievable on propagation through a slow-light medium. *Phys. Rev. A* **71,** 023801 (2 Feb. 2005).

a37. Moiseev, S. & Kröll, S. Complete reconstruction of the quantum state of a single-photon wave packet absorbed by a Doppler-broadened transition. *Physical review letters* **87,** 173601 (2001).

a38. Moiseev, S., Tarasov, V. & Ham, B. Quantum memory photon echo-like techniques in solids. *Journal of Optics B: Quantum and Semiclassical Optics* **5,** S497 (2003).

a39. Alexander, A., Longdell, J., Sellars, M. & Manson, N. Photon echoes produced by switching electric fields. *Physical review letters* **96,** 043602 (2006).

a40. Hétet, G. *et al.* Photon echoes generated by reversing magnetic field gradients in a rubidium vapor. *Opt. Lett.* **33,** 2323–2325 (2008).

a41. Sparkes, B. M. *et al.* Precision Spectral Manipulation: A Demonstration Using a Coherent Optical Memory. *Phys. Rev. X* **2,** 021011 (2 June 2012).

a42. Tittel, W. *et al.* Photon-echo quantum memory in solid state systems. *Laser and Photon. Rev.* **4,** 244–267 (2 2010).

a43. Afzelius, M., Simon, C., De Riedmatten, H. & Gisin, N. Multimode quantum memory based on atomic frequency combs. *Physical Review A* **79,** 052329 (2009).

a44. De Riedmatten, H., Afzelius, M., Staudt, M. U., Simon, C. & Gisin, N. A solid-state light–matter interface at the single-photon level. *Nature* **456,** 773–777 (2008).

a45. Afzelius, M. *et al.* Demonstration of Atomic Frequency Comb Memory for Light with Spin-Wave Storage. *Phys. Rev. Lett.* **104,** 040503 (4 Jan. 2010).

a46. Nunn, J. *et al.* Mapping broadband single-photon wave packets into an atomic memory. *Phys. Rev. A* **75,** 011401 (1 Jan. 2007).

a47. Riem, K. F. *et al.* Towards high-speed optical quantum memories. *N. Photon.* **4,** 218–221 (Mar. 2010).

a48. Reim, K. F. *et al.* Multipulse Addressing of a Raman Quantum Memory: Configurable Beam Splitting and Efficient Readout. *Phys. Rev. Lett.* **108,** 263602 (26 June 2012).

a49. Hosseini, M., Campbell, G., Sparkes, B. M., Lam, P. K. & Buchler, B. C. Unconditional room temperature quantum memory. *Nat. Phys., DOI: 10.1038/NPHYS2021.* doi:`DOI:10.1038/NPHYS2021` (2011).

a51. Geng, J. *et al.* Electromagnetically induced transparency and four-wave mixing in a cold atomic ensemble with large optical depth. *New Journal of Physics* **16,** 113053 (2014).

a53. Campbell, G., Hosseini, M., Sparkes, B. M., Lam, P. K. & Buchler, B. C. Time- and frequency-domain polariton interference. *New J. of Phys.* **14,** 033022 (2012).

a54. Campbell, G. T. *et al.* Configurable Unitary Transformations and Linear Logic Gates Using Quantum Memories. *Phys. Rev. Lett.* **113,** 063601 (6 Aug. 2014).

a57. Gorshkov, A. V., André, A., Lukin, M. D. & Sørensen, A. S. Photon storage in Λ-type optically dense atomic media. II. Free-space model. *Phys. Rev. A* **76,** 033805 (3 Sept. 2007).

a58. Hetet, G., Longdell, J., Alexander, A., Lam, P. K. & Sellars, M. Electro-optic quantum memory for light using two-level atoms. *Physical review letters* **100,** 023601 (2008).

a59. Dennis, G. R., Hope, J. J. & Johnsson, M. T. XMDS2: Fast, scalable simulation of coupled stochastic partial differential equations. *Computer Physics Communications* **184,** 201–208. ISSN: 0010-4655 (2013).

a60. Werner, M. & Drummond, P. Robust Algorithms for Solving Stochastic Partial Differential Equations. *Journal of Computational Physics* **132,** 312–326. ISSN: 0021-9991 (1997).

a61. Steck, D. A. Rubidium 85 D Line Data. *http://steck.us/alkalidata/rubidium85numbers.pdf* (2001).

a62. Steck, D. A. Rubidium 87 D Line Data. *http://steck.us/alkalidata/rubidium87numbers.pdf* (2001).

a63. Weller, L. *et al.* Cooperative Enhancement of Energy Transfer in a High-Density Thermal Vapor. eprint: `arXiv:1308.0129` (2013).

a64. Phillips, N. B., Gorshkov, A. V. & Novikova, I. Optimal light storage in atomic vapor. *Phys. Rev. A* **78,** 023801 (2 Aug. 2008).

a65. Lauk, N., O'Brien, C., Fleischhauer, M. & Brien, C. O. Fidelity of photon propagation in electromagnetically induced transparency in the presence of four-wave mixing. *Physical Review A* **88,** 013823. ISSN: 1050-2947 (July 2013).

a66. Phillips, N. B., Gorshkov, A. V. & Novikova, I. Slow light propagation and amplification via electromagnetically induced transparency and four-wave mixing in an optically dense atomic vapor. *Journal of Modern Optics* **56,** 1916–1925. ISSN: 0950-0340 (Oct. 2009).

a67. Phillips, N. B., Gorshkov, A. V. & Novikova, I. Light storage in an optically thick atomic ensemble under conditions of electromagnetically induced transparency and four-wave mixing. *Physical Review A* **83,** 063823. ISSN: 1050-2947 (June 2011).

a68. Camacho, R., Pack, M., Howell, J., Schweinsberg, A. & Boyd, R. Wide-Bandwidth, Tunable, Multiple-Pulse-Width Optical Delays Using Slow Light in Cesium Vapor. *Physical Review Letters* **98,** 153601. ISSN: 0031-9007 (Apr. 2007).

a69. Zhang, S., Zhou, S., Loy, M. M. T., Wong, G. K. L. & Du, S. Optical storage with electromagnetically induced transparency in a dense cold atomic ensemble. *Optics letters* **36,** 4530–2. ISSN: 1539-4794 (Dec. 2011).

a70. HAPPER, W. Optical Pumping. *Rev. Mod. Phys.* **44,** 169–249 (2 Apr. 1972).

a71. Nunn, J. *et al.* Multimode memories in atomic ensembles. *Physical Review Letters* **101,** 260502 (2008).

a72. Hush, M. R., Carvalho, A. R. R., Hedges, M. & James, M. R. Analysis of the operation of gradient echo memories using a quantum input-output model. *New Journal of Physics* **15,** 085020 (2013).

a73. Moiseev, S. A. & Arslanov, N. M. Efficiency and fidelity of photon-echo quantum memory in an atomic system with longitudinal inhomogeneous broadening. *Phys. Rev. A* **78,** 023803 (2 Aug. 2008).

a75. Lvovsky, A. I. Iterative maximum-likelihood reconstruction in quantum homodyne tomography. *Journal of Optics B: Quantum and Semiclassical Optics* **6,** S556 (2004).

a76. Roch, J. F. *et al. Phys. Rev. Lett.* **78,** 634 (1997).

a77. Ralph, T. C. & Lam, P. K. Teleportation with Bright Squeezed Light. *Phys. Rev. Lett.* **81,** 5668 (1998).

a78. Hetet, G., Peng, A., Johnsson, M., Hope, J. & Lam, P. K. Characterization of electromagnetically-induced-transparency-based continuous-variable quantum memories. *Physical Review A* **77,** 012323 (2008).

a79. Bowen, W. P. *et al.* Experimental investigation of continuous-variable quantum teleportation. *Phys. Rev. A* **67,** 032302 (2003).

a80. Higginbottom, D. B. *et al.* Spatial-mode storage in a gradient-echo memory. *Phys. Rev. A* **86,** 023801 (2 Aug. 2012).

a81. Appel, J., Figueroa, E. & Lvovsky, A. Adiabatic frequency conversion of optical information in atomic vapor. *Opt. lett.* **32,** 2771–2773 (2007).

a82. Campbell, G., Ordog, A. & Lvovsky, A. I. Multimode electromagnetically induced transparency on a single atomic line. *New J. of Phys.* **11,** 106021 (2009).

a83. Reck, M., Zeilinger, A., Bernstein, H. J. & Bertani, P. Experimental realization of any discrete unitary operator. *Phys. Rev. Lett.* **73,** 58–61 (1 July 1994).

a84. Hetet, G., Longdell, J. J., Sellars, M. J., Lam, P. K. & Buchler, B. C. Multimodal Properties and Dynamics of Gradient Echo Quantum Memory. *Phys. Rev. Lett.* **101,** 203601–4 (2008).

a85. Bimbard, E., Jain, N., MacRae, A. & Lvovsky1, A. I. Quantum-optical state engineering up to the two-photon level. *Nat. Phot.* **4,** 243–247 (2010).

a86. Tittel, W., Brendel, J., Zbinden, H. & Gisin, N. Quantum Cryptography Using Entangled Photons in Energy-Time Bell States. *Phys. Rev. Lett.* **84,** 4737–4740 (May 2000).

a87. Andersen, U. L., Filip, R., Fiurášek, J., Josse, V. & Leuchs, G. Experimental purification of coherent states. *Phys. Rev. A* **72,** 060301(R) (2005).

a88. Menicucci, N. C., Flammia, S. T. & Pfister, O. One-Way Quantum Computing in the Optical Frequency Comb. *Phys. Rev. Lett.* **101,** 130501 (13 Sept. 2008).

a89. Pysher, M., Miwa, Y., Shahrokhshahi, R., Bloomer, R. & Pfister, O. Parallel Generation of Quadripartite Cluster Entanglement in the Optical Frequency Comb. *Phys. Rev. Lett.* **107,** 030505 (3 July 2011).

a90. Goldschmidt, E. A., Beavan, S. E., Polyakov, S. V., Migdall, A. L. & Sellars, M. J. Storage and retrieval of collective excitations on a long-lived spin transition in a rare-earth ion-doped crystal. *Opt. Express* **21,** 10087–10094 (Apr. 2013).

Books

b55. Lambropoulos, P. & Petrosyan, D. *Fundamentals of Quantum Optics and Quantum Information* ISBN: 978-3-540-34571-8 (Springer Berlin Heidelberg NewYork, 2007).

b56. Townsend, J. S. *A Modern Approach to Quantum Mechanics* ISBN: 978-81-309-1314-8 (University Science Books, 2010).

b74. Leonhardt, U. *Essential Quantum Optics: From Quantum Measurements to Black Holes* ISBN: 9780521869782 (Cambridge University Press, 2010).

Other resources

c50. Hosseini, M. *Quantum Optical Storage and Processing Using Raman Gradient Echo Memory* PhD thesis (The Australian National Univeristy, 2012). <`http://photonics.anu.edu.au/theses/2012hosseini.pdf`>.

c52. Sparkes, B. M. *Storage and Manipulation of Optical Information Using Gradient Echo Memory in Warm Vapours and Cold Ensembles* PhD thesis (The Australian National Univeristy, 2013). <`http://photonics.anu.edu.au/theses/2013Sparkes.pdf`>.

II

Monolithic resonators

CHAPTER 12

Introduction

In this part of the thesis I will present results on the fabrication of monolithic resonators made from optical crystals. The goal of the work is to develop high-finesse crystalline resonators for applications in quantum optics. The resonators confine light to a small mode volume while maintaining long photon lifetimes, making them an excellent platform for large optical non-linearities. We produced resonators in two geometries: whispering gallery resonators and square resonators that use total internal reflection to confine the optical modes. Both geometries provide bulk absorption limited linewidths in $LiNbO_3$, fast voltage tuning and variable coupling to free-space modes using prism couplers.

12.1 Overview of monolithic resonators

Optical resonators are a central building block for studies that involve coherent optics. They consist of a set of reflecting boundary conditions that define a closed-loop path for light such that it interferes with itself on subsequent round trips along the path. As discussed in section 3.5 this reduces the mode density to only modes that undergo constructive interference. It is this restriction of the available optical modes that enable single-mode lasers which are an enabling technology for fibre telecommunications and are the workhorse for studies in laser spectroscopy, metrology, advanced microscopy, optical quantum science and many other fields.

Given the broad scope of applications for optical resonators it is unsurprising that there is an equally broad range of physical designs for resonators. These range from the bulk-optic resonators that form the lasing cavities of the Ti:Sapphire lasers used in the quantum memory experiments of this thesis, to the dielectric coatings used for every reflective and anti-reflective surface on the various lenses, mirrors, and beam-splitters. There is also significant interest in developing new types of resonator, particularly compact designs [b1] that could be used for miniaturising lasers and other optical elements such that they can fit into modular equipment or even directly onto integrated optical circuits. Most small resonators, from wavelength-scale up to a few millimetres, are monolithic, a term that refers to a structure that is composed of only a single piece of material. These structures do not use separate reflective optics to confine light to an enclosed region but rather are themselves structured to provide the appropriate boundary conditions.

Examples of monolithic resonators include whispering gallery resonators, micro-toroids, and photonic-crystal resonators, all of which have attracted considerable research interest in recent years. More conventional types of monolithic resonator may also include resonators that have separate reflective components, such as dielectric coatings or even entire mirrors, that are distinct from the solid substrate but are adhered to it permanently such that the entire structure acts mechanically as a solid block. With a sufficiently broad definition of monolithic, even the mode-cleaner cavities of the experiments in this thesis could qualify because the structural element of the cavity is a single machined block of aluminium or Invar.

The monolithic resonators that are the subject of this chapter cover a small subset of possible designs. They are millimetre-scale resonators that each consist of a single piece of transparent optical material that can guide light along a closed path within the material itself. Light can either be continuously guided along the inner surface of the material [a2, a3] or can be reflected off mirrors formed by the interface between the resonator and the surrounding free-space. In the case of reflections off mirror surfaces, the mirrors can use dielectric coatings [a4], total internal reflection [a5–a7], or both [a8] to reflect the light.

Monolithic resonators that use dielectric coatings have the advantage of a simple coupling interface to free-space. The coatings can be engineered to transmit a determined ratio of incident light for a given wavelength. In contrast, total internal reflection resonators need to be evanescently coupled, which provides broadband operation [a7] and the ability to tune the coupling strength [a5] at the expense of added complexity.

Before beginning to cover the original work that I have done, I will begin with a brief literature review on just a few of the applications of whispering gallery resonators and crystalline monolithic resonators. This is only a small subset of the work being done on these devices and the crystalline resonators themselves form a small subset of the work done on monolithic resonators in general.

12.2 Whispering gallery resonators

One of the most diverse types of compact monolithic resonator is the whispering gallery resonator. These take their name from the whispering gallery in St. Paul's Cathedral in London. The building features a circular wall that guides acoustic waves along the inside perimeter, allowing a listener close to the wall to hear whispers spoken anywhere else around the wall's perimeter. The guided acoustic waves are called whispering gallery waves, and lend their name to optical whispering gallery modes formed by circular reflecting boundary conditions.

Optical whispering gallery resonators use total internal reflection to guide light around a toroidal surface. The resonator is therefore made from a material with a higher index of refraction than its surroundings, and can be almost any dielectric material for a resonator in the air. The circular geometry and use of total internal reflection makes whispering gallery resonators ideal for making compact, high-finesse optical cavities. The circular shape can be easily produced on a precision lathe for crystalline materials or simply by melting an amorphous material and allowing surface tension to pull the surface into a smooth radius. The technique of melting an amorphous material to form a resonator is known as surface tension induced re-flow.

At the small end of the size scale, resonators can be fabricated with diameters of only a few optical wavelengths. These are formed by using photo-lithography to produce small disks of silica that are raised off of a substrate [a9] or in tapered optical fibres [a10, a11]. An intense laser can then be used to melt and re-flow the perimeter of the disk, forming it into a smooth surface that can guide light. Larger versions of this type of structure can support optical modes with a lifetime limited only by material absorption [a3, a12, a13]. The extremely narrow line-widths that correspond to the long optical lifetimes make these resonators ideal for use as high-sensitivity sensors . Any environmental condition that can be coupled to the refractive index of the resonator, temperature for example, can be detected with a high precision. The evanescent field can also be used to couple to objects that are near the resonators, such as nearby atoms [a14] or particles [a15], due to the perturbation that they induce on the optical mode.

The high finesse, eq. (4.15), and small mode-volume also provides pathways to study opto-mechanics [a2, a16]. The circulating optical mode can be coupled to travelling acoustic waves in the resonators material [a17], bulk acoustic modes of the resonator itself [a18, a19], or even surface waves in fluids on the resonator surface [c20]. The small mode-volume is also of interest to non-linear optics, providing string non-linearities at low light levels. The next section will only discuss $\chi^{(2)}$ non-linearities [a13, a21] in crystalline materials but it is worth noting that $\chi^{(3)}$ non-linearities in amorphous materials have been used to demonstrate extremely compact devices for generating frequency combs [a22, a23].

12.3 Crystalline monolithic resonators

Crystalline monolithic resonators include whispering gallery resonators as well as other geometries. The technique of laser-induced re-flow cannot be used to produce them due to the crystalline structure. Instead, conventional optical fabrication techniques are used, such as grinding, machining, and polishing, all of which will be discussed in chapter 13.

Some of the earlier crystalline monolithic resonators were produced from $LiNbO_3$ in pursuit of exactly the same goals as this work [a6, a7]. These were fabricated using standard polishing techniques to produce square or triangular resonators with one or more spherical surfaces to form a stable resonator cavity. A variation is to produce a simple linear cavity with two spherical mirrors and use dielectric reflection coatings rather than total internal reflection. Such cavities have been used successfully to produce squeezed light.

More recent work on crystalline resonators has tended to focus on whispering gallery geometries due to the potential for a very high finesse if the polishing is sufficient and bulk absorption is low in the material that is used, such as with CaF_2 . A record quality factor eq. (4.17) of 10^{11} was obtained in such a resonator, corresponding to the longest photon lifetime recorded in a monolithic resonator of 0.5 ms [a12]. This type of resonator has also been used to generate frequency combs that span a large bandwidth thanks to the high non-linearity afforded by the low mode volume and high finesse.

Whispering gallery resonators that are made from electro-optic materials, usually $LiNbO_3$, have found application as both electro-optic devices and non-linear devices. The refractive index of these resonators can be tuned by applying an electric voltage across the resonator. This can be used as a fast electro-optic switch by bringing the resonator in or out of resonance such that it can transmit light from one coupling interface to another. They have also been used as high-order optical filters by chaining together multiple tunable resonators.

Research into developing $LiNbO_3$ whispering gallery resonators for non-linear and quantum optical process is also under-way. The high finesse has been used to achieve large $\chi^{(2)}$ non-linearities for the generation of photon pairs [a13, a21]. The small mode-volume afforded by whispering gallery resonators is also being investigated for the purpose of developing non-linear optical gates [a24].

12.4 Context of this work

While crystalline whispering gallery resonators are being actively pursued as a means to generate quantum light, so far, no demonstration of degenerate squeezing has been reported. These compact designs present significant challenges but the potential advantages of miniaturising quantum sources is sufficient incentive to overcome them.

CHAPTER 13

Fabrication

Crystalline whispering gallery resonators are fabricated on a diamond turning lathe. Diamond turning combines a number of technologies to produce high-precision surfaces by removing material from a bulk substrate. High-resolution optical position encoders, extreme rigidity bearings and finely-tuned motion controllers are used to control the location of a diamond cutting tool to nanometre precision. This motion accuracy, combined with the extreme hardness of diamond, enable the cutting edge of the tool to be moved across a work-piece in a path that deviates from the intended shape by only a fraction of an optical wavelength, enabling the direct machining of optics.

The high rigidity of the machining system, including the bearings, feedback motors, and general construction of the enclosure, also allows the machining of materials that are far too hard to be machined by traditional manufacturing process. These materials include crystals that are frequently used in optics, such as calcium fluoride, potassium titanyl phosphate (KTP), zinc selenide, silicon, germanium and lithium niobate. The most common method of diamond machining is diamond turning, in which a precision lathe is used to position a diamond cutter in relation to a rotating work-piece.

The ultra-precision lathe, a UPL250 supplied by Moore Nanotech [c25], has four axes of motion. The work-spindle is an air-bearing rotary axis that can rotate at up to 10,000 revolutions per minute or provide precise angular positioning synchronised to the other axes. Two axes provide linear motion, one along the axis of the work-spindle and the other perpendicular to it. An additional rotary axis provides precise rotation about an axis normal to the plane formed by the two linear axes and perpendicular to the work-spindle. Appendix C provides further details about the lathe.

Whispering gallery resonators are particularly well-suited to fabrication by diamond turning due to their cylindrical symmetry. The procedure for producing the resonators consists of three steps, although the details of each depend on the type of resonator that is being fabricated.

First a circular cut-out is produced from the bulk material, usually by grinding, shaped and mounted on a suitable post that can be attached to the work spindle. The resonator is then roughly shaped by relatively rapid diamond turning cuts and finished using shallow, slow cuts to produce an optical surface. The final step is to polish the resonator to achieve an extremely smooth surface for high quality factor resonators. Non-cylindrically symmetric resonators are fabricated using a similar process although the cutting is done at a lower spindle speed to allow the tool to move synchronously with the spindle to shape

Figure 13.1: A collection of fabricated whispering gallery resonators.

the part.

13.1 Preparing the blanks

The bulk crystal that is to be sculpted into resonators arrives from the manufacturer in one of three forms: rods, wafers or optic blanks. For large CaF_2 resonators, optic blanks with a $1/2''$ or $1''$ diameter can be purchased and machined with no further shaping. For other materials and smaller diameters, however, the bulk material must be cut using a grinding process.

Rods of crystal can be diced to form thin disks using a diamond wafering blade. These blades have an edge of resin embedded with diamond particles which can quickly grind through bulk material, leaving a flat but rough edge due to non-ductile removal of material. The crystal rods are mounted and centred on the lathe's work spindle and the the wafering blade is mounted on a second air-bearing grinding spindle which can be mounted on the rotary axis of the lathe.

The rotary axis is used to set the wafering blade to be in the plane perpendicular to the axis of the crystal rod and the blade is used to slice into the side of the rod by moving the linear axes of the lathe. This allows the grinding to be done with all the precision of the lathe's linear axes. The rod is rotated slowly on the work spindle while the rapidly rotating wafering blade is cutting through it to ensure a cylindrically symmetric cut. A generous quantity of coolant mist is applied to the blade and the rod, as is show in fig. 13.2. A small nub is often left on the cut disk when it is separated from the bulk rod but this is easily removed using the lathe.

A similar procedure is used to extract circular cut-outs from thin crystal wafers. The

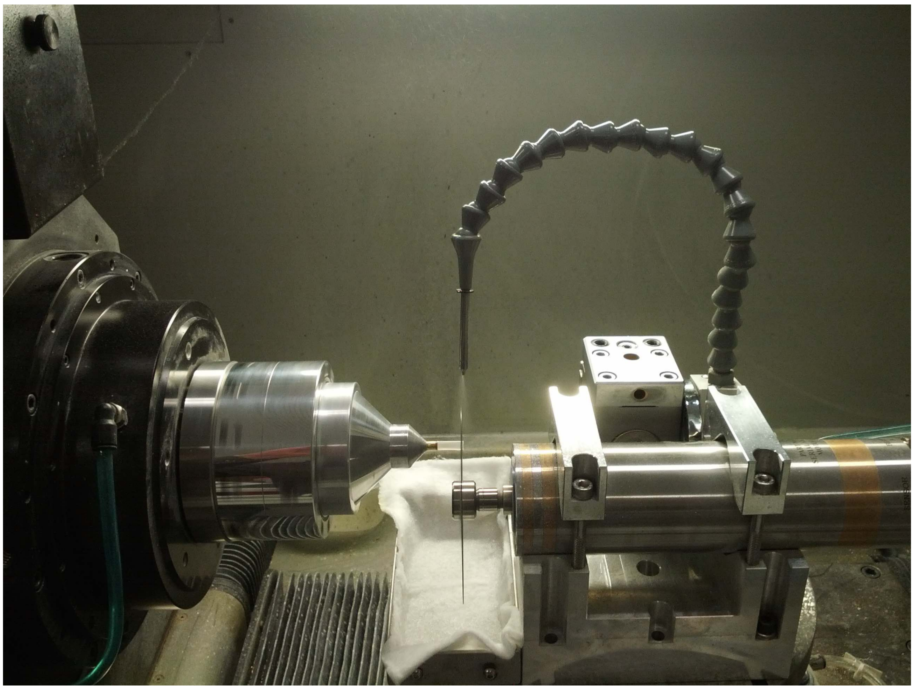

Figure 13.2: A rod of $LiNbO_3$ is diced into disks that will be machined into monolithic resonators.

wafers are bonded to a flat piece of aluminium using cyanoacrylate adhesive (super-glue) which is then mounted to the work spindle. A hollow coring bit, again using a resin-bonded diamond abrasive cutting edge, is mounted to the grinding spindle and is advanced into the wafer. Here the spindle is held stationary and is used only to position the correct part of the wafer under the bit. The aluminium mounting is easily cut by the diamond grinding tip so the wafer can be cut through entirely.

These procedures for making disks are sufficient for producing small numbers of disks but have some drawbacks. The produced disks occasionally get blown away by coolant and lost so care must be taken to collect them. Additionally, when extracting cores from wafers, the stresses from cutting one core from the wafer tend to fracture some of the surrounding material, reducing the number of disks that can be made.

Once the crystals are roughly shaped as disks they can be mounted to a post or a stud that can then be attached to the lathe. Large resonators are simply glued to a $1/2''$ optic post which can then be inserted into a holder. This has the advantage of being very straightforward to position in an experiment, being compatible with standard post-holders, however care must be taken to avoid touching the diamond tool to the mounting post while machining because standard steel posts cannot be machined with diamond tips. Figure 15.3 shows a CaF_2 resonator glued to a standard optic post using epoxy.

For smaller resonators or resonators that need to have conductive plates immediately next to the optical path steel posts are not an option because the post needs to be machinable. For these, we fix the resonator blanks to a small brass stud that features a through-hole and an M3 thread. This allows the stud to be screwed into a conical holder, shown in fig. 13.3, that serves to mount the stud to the lathe's work spindle via the vacuum chuck.

The brass stud can be machined with the same diamond tools that are used for the crystal. This allows the cut path to move continuously through brass and crystal such that the resonator has nearly the same diameter as the stud. This provides the largest surface for the adhesive to stick to, mechanical support (and damping) for the resonator during

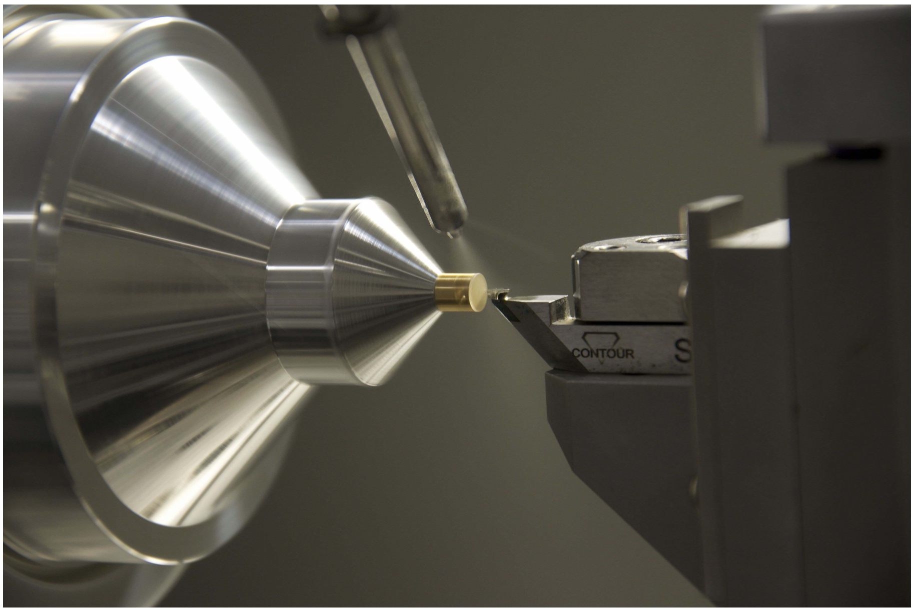

Figure 13.3: Flatting the face of a $LiNbO_3$ resonator blank. The blank is epoxied to a brass stud which is, in turn, screwed in to an aluminium holder.

cutting and a uniform contact for thermal conductivity or applied electric fields.

The cutting forces while machining brittle crystal have, counter-intuitively, been found to be roughly independent of the crystal radius for a fixed spindle speed and feed-rate [a26]. The area of contact between the resonator and mounting stud, however, is quadratic in radius. This places a limit on the minimum size of the resonator simply due to the strength of the bond to the stud. For this reason, the type of adhesive used is important to the production of small resonators as well as the surface finish of the resonator and the stud.

A poor surface finish on the surface resonator produces a stronger adhesive bond than a smooth one. To improve the bond with the stud, the surface of the stud was machined on the lathe with channels cut into its surface using a non-cylindrically symmetric cut profile appendix C.2 shown in fig. 13.4. The channels provide a path for the epoxy to be squeezed out for between the resonator and the stud, reducing air bubbles, providing a more uniform thickness of epoxy and providing more structure for the epoxy to adhere to. To further improve the bond quality, chemical abrasion of the brass mounting stud using ferric chloride was also tested but did not significantly improve the failure rate of the bond.

A variety of epoxies were used, including Araldite 2012, 2014, and 2021, and various brands of consumer grade adhesives. We did not conduct any systematic study of the bond strength for the resonators but found that the Araldite 2014 and 2021 seemed to have good bond strengths. Consequently, we favoured them over the other epoxies due to the improved chemical and temperature resistance. These properties were found to be important because the resonators were subjected to phase-matching temperatures of over 60° C and cleaned using acetone and methanol. In at least one case, the epoxy failed and the resonator fell off its mounting stud after being repeatedly polished, cleaned and held at the phase-matching temperature.

Once the circular crystal blank was mounted to the stud, it was machined into approximately correct dimensions using roughing cuts. These are performed using tools that have become too dull to provide a good surface quality and are usually much deeper than

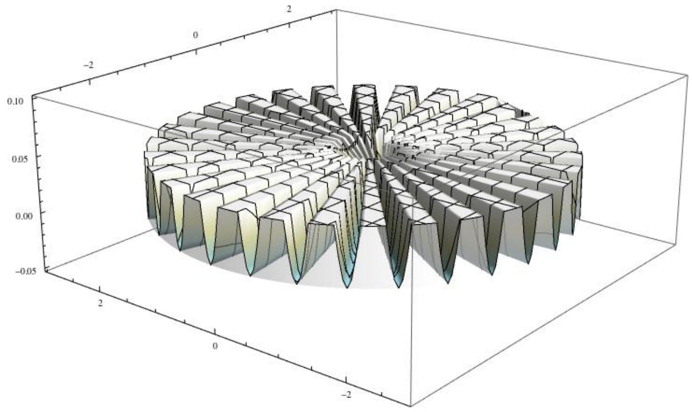

Figure 13.4: Channels are cut into the mounting stud to facilitate an even distribution of epoxy.

finishing cuts to increase the material removal rate.

13.2 Cutting

The resonator is machined by attaching the substrate to the work spindle of the high precision lathe and spinning it with constant velocity while the two linear axes of the lathe act to shape the profile. With this configuration, cuts can be made to be exceptionally shallow by maintaining a high spindle velocity while slowly moving the diamond tool along the resonator substrate, forming a shallow, helical cut. The goal of the machining process is to force the material being machined to deform in a ductile manner while being cut.

The particular cutting parameters required to maintain a ductile cut depend on the material, as well as the crystallographic orientation of the material. The dependence on the crystal orientation is a particular challenge for producing whispering gallery resonators because it is constantly changing as the work-piece rotates under the diamond tool. An added complication is that hard optical crystals will rapidly dull the diamond cutting tool, changing the parameters that will produce a ductile mode cut from one pass to the next.

While the cutting parameters serve as a rough guide of how to obtain ductile-mode cutting, the dependence of the cut quality on the sharpness of the tool tends to be the most important factor. As the tool dulls, finishing cuts can be made increasingly shallow to maintain a finishing cut; however, in hard materials ($LiNbO_3$) this may only extend the finishing lifetime of the tool by only a few passes. The best strategy is to keep track of which portions of a tool's cutting edge have been used and to attempt to equalise wear along the entire edge.

The rate of tool wear depends strongly on the material that is being cut as well as the type of cut. A new tool may be capable of only a handful of finishing passes in $LiNbO_3$ but hundreds in CaF_2. Furthermore, a tool that can no longer provide a finishing cut in a hard material may still provide an adequate surface in softer material. Once a tool is no longer sharp enough to perform completely ductile-mode material removal it can continue to be used as a roughing tool, thereby saving wear on still-sharp tools.

It is sometimes simpler to abandon attempting to obtain a perfect ductile cut and to simply polish away defects. This is particularly true for CaF_2 and MgF_2 which tend to develop large chips in the event of a brittle fracture. These two materials are also relatively soft which, combined with the cylindrical symmetry of the whispering gallery resonators, makes them straightforward to polish.

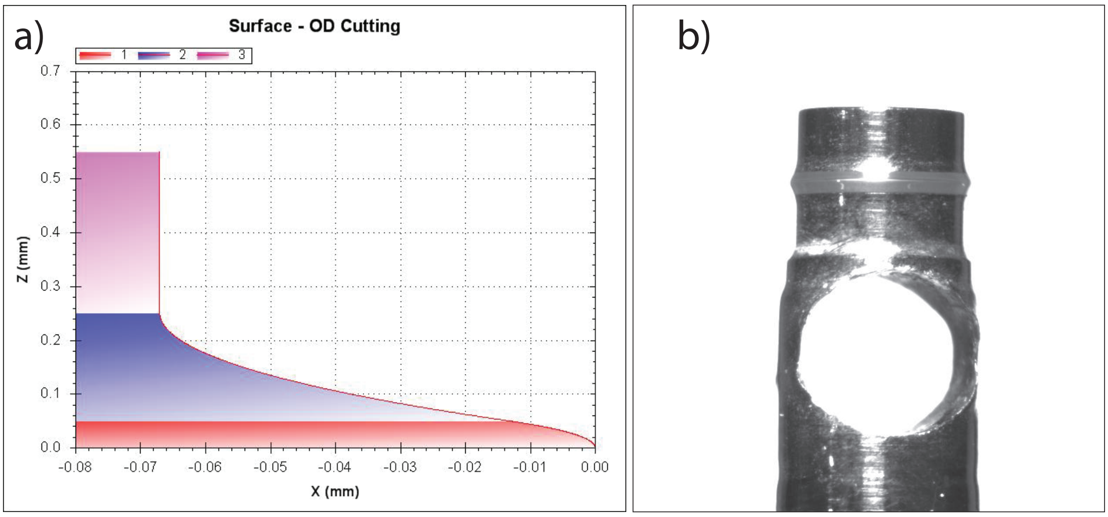

Figure 13.5: a) Programmed tool path for a 1 mm minor radius whispering gallery resonator. b) An early example of a finished $LiNbO_3$ whispering gallery resonator mounted on a brass stud and capped with a brass electrode. Future iterations would feature a tapered stud with a lower through-hole to avoid the thin walls of the stud shown here.

The numerical control (NC) program for a whispering gallery resonator is simple to produce. It consists of a list of points, defined in cylindrical coordinates, that provide the radius and height at every point along the resonator. These can be produced by software (proprietary of Moore Nanotech) that joins arc and line segments into a continuous path. An example is shown in fig. 13.5. Concave blend arcs are added to either side of the convex resonator profile to ensure that there exists in the cut no concave radius of curvature that is less than the radius of the cutter itself. If this were not done, a large portion of the cutter edge would briefly be in contact with the material, causing a large increase in cutting forces.

13.3 Ductile mode machining of brittle materials

The high stiffness that is an integral part of diamond turning machines allows materials that are typically too brittle for machining to be shaped by diamond turning. The key to achieving an acceptable surface quality is maintain the cut parameters such that material is to removed as a continuous sheet, rather than chipping it out in chunks. In metals, this is the natural mode of removal due to their high ductility. Brittle materials, however, are prone to fracture under stress. Machining of these requires the cutting parameters to be carefully tuned to induce appropriate stresses on the crystal structure that result in continuous material removal rather than chipping [a27].

A number of studies have been performed on a variety of materials to model this mode of material removal and find suitable cutting parameters. While these studies have found various (and often contradictory) regimes of parameters for each material examined, a number of reasonably universal conclusions can be drawn: the chip thickness should be small compared to cuts in metal, the tool should have a negative rake and a moderate amount of coolant should be used.

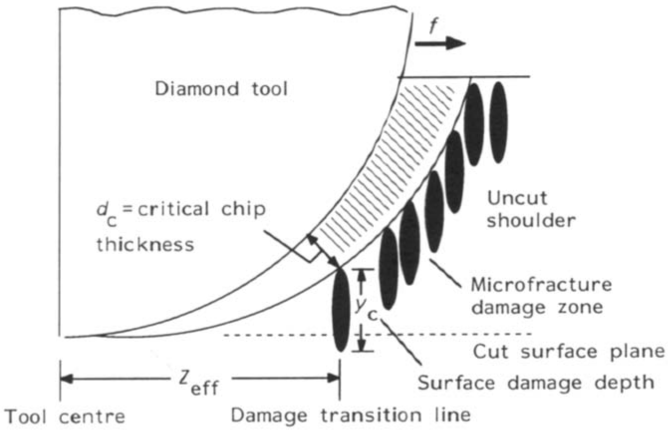

Figure 13.6: The removed chip thickness for one pass of a round diamond cutter through a substrate varies along the edge of the cutter [a28]. Below a critical thickness, material is removed in a ductile mode, above this thickness fractures form.

13.3.1 Chip thickness

The depth of the cut is an obvious parameter that needs to be controlled. If the chip is being machined is too thick, the material will fracture under the cutting forces or, in the most extreme cases, be flung from the work spindle. Figure 13.6, reproduced from [a28], shows the shape of the removed chip if the cutting is performed with a round cutter. The goal is to tune the feed-rate per revolution and the cut depth so that any fractures generated by cutting forces are removed by the ductile mode region of the on subsequent passes of the tool.

A number of studies have investigated the critical chip thickness for many materials under a variety of cutting conditions [a29]. In particular, much research has been done on cutting silicon [a30–a32], germanium [a28, a33], lithium niobate [a34] and calcium fluoride [a35, a36].

Once the critical chip thickness has been determined for a particular material and a set of machining conditions, it should be possible to calculate the feed-rate and cut depth that can be used for a tool of a given radius. In practice, this is used as a rough guide; however, other factors such as tool wear were found to play such an important role that no single number could be placed on a chip thickness for any particular material.

In general, smaller tool radii necessitate a slower cut to guarantee that fractures do not form during cutting. A slower or shallower cut can also be used to offset some of the effects of tool wear, yielding good finishes even as the edge of the cutter becomes worn by reducing the chip thickness. A slower cut, however, results in more cutting time per part, increasing the rate at which the tool is worn down.

13.3.2 Rake angle

The rake angle is the angle formed between the cutting edge of the tool and the part being cut, relative to 90°. A positive rake angle corresponds to an acute angle and a negative rake angle corresponds to an obtuse angle. For example, a rake angle of 10° positive means an 80° angle between the work-piece and the top face of the cutting tool. The rake angle can be either built into the tool or adjusted by altering how it is mounted on the lathe.

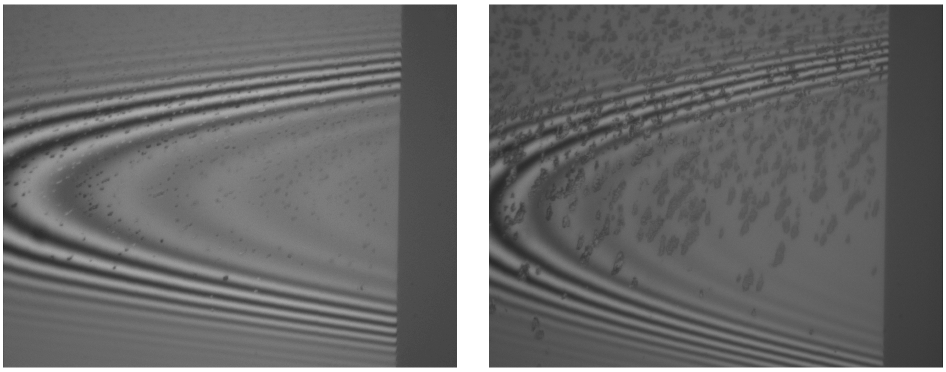

Figure 13.7: Two finishing cuts made with the same cutting parameters and the same diamond tool. The surface shown on the left was cut first and shows some surface damage from fracturing. The cut was repeated four more times and the surface was imaged again. Increased surface damage due to the tool wear incurred during the additional cuts is clearly visible.

A −25° rake angle is commonly set into tools which are sold for the purpose of machining brittle materials. This is suitable for use in a variety of crystals but has the disadvantage that the tool can only be used in one orientation; rotating the tool to use a different part of the cutting edge changes the rake angle. Another type of tool, the “special edge" tool, uses a microscopic rake at the cutting edge of an otherwise flat (0° rake) cutter. These special edge tools largely avoid the problem of finding an optimal set of cutting parameters: we found that a sharp special edge tool will provide a ductile mode cut in both lithium niobate and calcium fluoride over a much larger range of cutting parameters than is found in literature with standard tools [a34, a35].

13.3.3 Tool wear

Diamond is the hardest bulk material but, even so, it is not impervious to wear. Cutting hard optical crystals, or even soft metals, will dull the cutting edge of a diamond tool. This wear has a detrimental effect on the tools ability to perform ductile mode cuts and will eventually even affect the form accuracy of a part due to deviation of the cutter from a circular profile on the scale of hundreds of nanometres.

There have been extensive studies on tool wear in diamond turning [a37, a38]; however, our studies have been primarily qualitative in nature. We found that tool wear is the most important factor in achieving a ductile cut, but it depends heavily on the type of cuts being made. Remarkably aggressive cuts in lithium niobate, using similar parameters to what would be suitable for finishing passes in malleable metals, result in ductile cuts if the tool is completely new. A worn tool, on the other hand, will produce ductile cuts only for a small and changing set of cutting parameters, if at all. An example of the effect of tool wear is shown in section 13.3.3. Five identical cuts were made in $LiNbO_3$ and images were taken after the first and last cut, clearly showing more severe surface damage after the final cut.

The tool lifetime depends on the type of material being cut and the required surface quality. Each tool can only cut a few tens of millimetres in lithium niobate before fractures occur in subsequent cuts. The exact number depends very much on the radius of the tool and the cutting parameters which necessitates manual monitoring of the cut quality. Each set of finishing cuts on lithium niobate should be performed with a freshly re-lapped tool but calcium fluoride is more forgiving, allowing a handful of parts to be produced by each tool. This description is deliberately vague due to the large variability of the lifetime.

13.3.4 Crystal orientation

One of the key factors governing the quality of a cut in a crystalline material is the crystal orientation: cuts along certain crystallographic planes are more susceptible to fracture. As the work-piece rotates under the diamond tool, the cutting forces are directed at constantly changing angles with respect to the crystalline structure. In general, crystals tend to cleave along planes that lies parallel to planes of high density.

A set of cutting parameters that works well for certain orientations may cause fractures at different orientations. In some cases, I have observed a reduction in chip thickness to improve the cut quality at one orientation while causing another orientation to develop fractures which were not present with the more aggressive cut. Generally, however, a smaller chip thickness will improve the quality at all orientations and parameters which produce good cuts at the orientations which are most susceptible to fracture will also produce good cuts at the more resilient orientations.

13.4 Lithium Niobate

Our earliest attempts at machining optical crystals were done with $LiNbO_3$. After many hours of sweeping various quantities such as spindle speed, cut depth, feed-rate and tool radius we had little in the way of reproducible results. Ultimately it was found that the tool wear in $LiNbO_3$ is quite severe and changes the quality of a cut significantly even if all other parameters are held constant. Once this was discovered, machining $LiNbO_3$ became straightforward; fractures tend to be small and predictably located and the material is robust against accidental damage (within reason).

The $LiNbO_3$ crystal has three-fold rotational symmetry, see appendix B.1.1, which results in a periodic pattern in the cut quality. Each 120° section features two regions which are susceptible to fracture, as shown in fig. 13.8. This is likely due to cleavage along the $\{10\bar{1}0\}$ and $\{11\bar{2}0\}$ lattice planes, although we would need to more carefully orient the crystal about its c-axis to verify this. The natural cleavage plane, $\{01\bar{1}2\}$ is not cut directly (due its orientation of forming a 32.75° angle with respect to the c axis) but may be a factor in the size and shape of fractures. These regions will feature damage that reduces in severity as smaller cut depths are used and can be eliminated by the use of a finishing cut with a sharp tool.

Very dull tools will produce noise while cutting $LiNbO_3$, likely due to acoustic modes being excited in the crystal. This almost guarantees that the cut quality will be poor and increases the possibility of causing a large fracture through the work-piece.

Roughing cuts can be very aggressive in $LiNbO_3$. Even a deep cut with a dull tool will usually cause only superficial damage to crystal. A cut depth of 10 μm with a feed-rate of 2.5 μm/rev is a typical rough cut in $LiNbO_3$ using a moderately worn diamond tool with a 0.2 mm radius. These are the parameters used for the cut shown in fig. 13.8. The largest fractures that result are usually a few tens of microns across and are not usually deeper than 15 microns. A deeper or faster cut begins to risk putting enough strain on the crystal that fractures will propagate much further into the substrate. For comparison a rough cut in metal is typically done with a 10 micron depth of cut and a 10 μm/rev feed-rate.

Finishing cuts must be done with a diamond cutter which has a fresh cutting edge to obtain a ductile cut at all crystal orientations. The diamond tools are rapidly worn by the

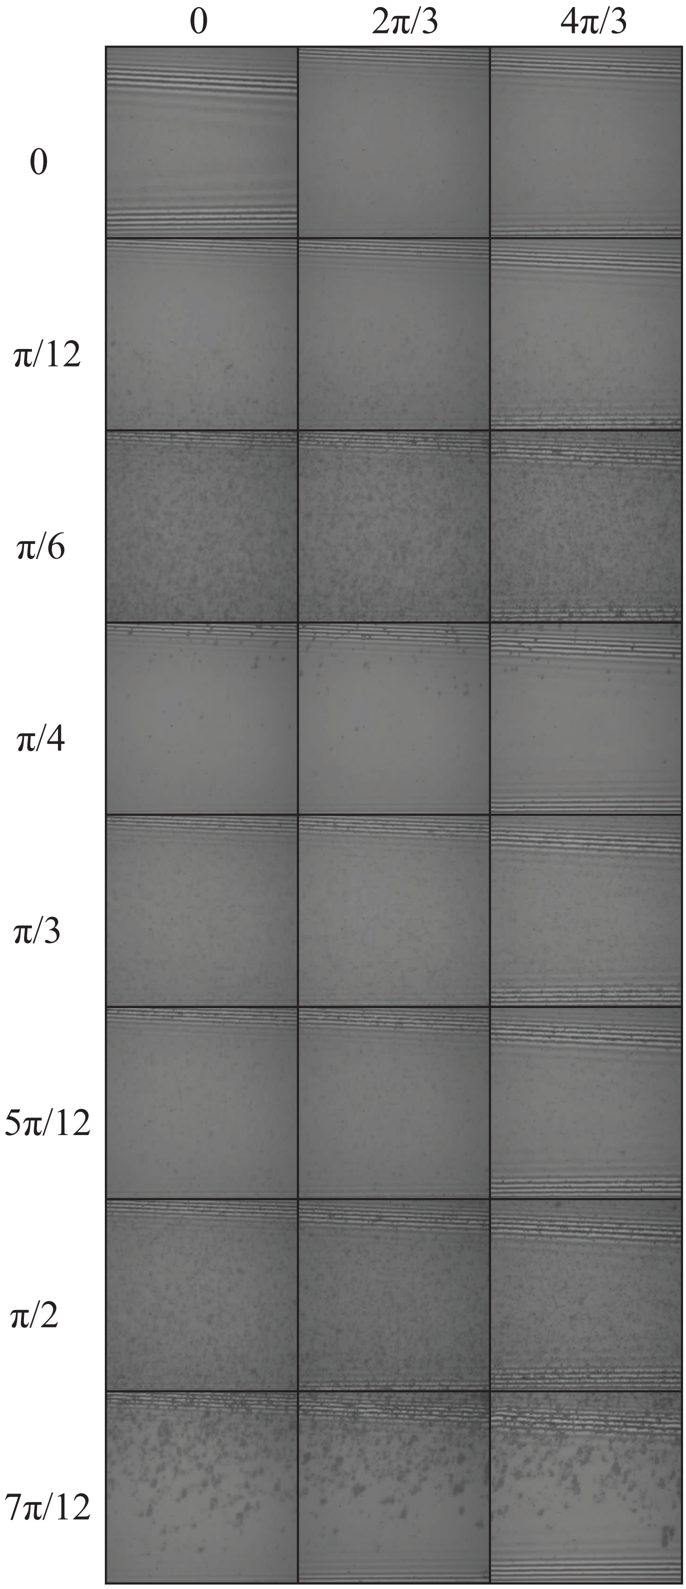

Figure 13.8: Images of a roughing cut in $LiNbO_3$ taken every 15° degrees using an on-machine microscope. The three-fold rotational symmetry of the crystal structure is illustrated by showing each 120° segment beside the preceding one. Interference fringes can be seen at the top and bottom of each image due to the interferometry objective used.

process of cutting $LiNbO_3$ and the lifetime for a finishing cut can be as short as a few passes over a work-piece. Section 13.3.3 shows the effect of tool wear: two identical cuts were performed with a tool that was beginning to show signs of damage, separated by four passes. The increased damage caused by wear incurred during the intervening passes is clear: the right image, taken after the additional cuts shows larger fractures than the earlier cut.

A typical finishing cut for a 0.5 mm radius cutter would be a 3 μm depth of cut and a feed-rate of 1 μ/rev. This produces a surface with minimal residual roughness due to the shape of the tool; however, even a cut depth of 5 μm and a feed-rate of 3 μ/rev will produce a ductile mode cut if the tool is new.

13.5 Calcium fluoride

Calcium fluoride is a softer material than lithium niobate and, consequently, tool wear occurs more slowly. However, brittle fractures in CaF_2 tend to produce produce larger defects. In general, the roughing passes must be relatively shallow, one to three μm, to avoid producing surface damage that extends up to 20 μm into the material. Additionally, the material is more prone to occasional chips in otherwise ductile cuts. If these occur at the apex of a whispering gallery mode resonator, further cuts must must be made to remove them. Careful monitoring of the surface after each finishing pass is therefore important to produce reliable finishes.

The relative softness of the material results in short times required for polishing. Tens of microns of material can be removed relatively quickly by polishing so for applications that do not require a precise form accuracy brittle fractures may be corrected by polishing rather than attempting numerous cuts.

13.6 Polishing

After the whispering gallery resonators are shaped and given a finishing cut, they can be polished to remove remaining surface roughness. This is done using a polishing slurry that consists of microscopic diamond particles suspended in a lubricant. The slurry is then applied to a soft fabric which is used to polish the resonator.

Polishing is process whereby material is removed from a surface by abrasion by small, hard particles that are carried by a soft, elastic pad or by a fluid jet. The small scale of the cutting edges, microscopic granules of diamond or a hard ceramic, results in ductile removal of even brittle materials [a27], ensuring that no fractures develop in the work-piece. The rate of material removal and the resulting smoothness of the surface finish are determined by the properties of the polishing pad as well as the size and type of polishing particles.

Whispering gallery resonators can be quickly polished by leaving them on the work spindle after cutting and using the lathe to rotate them. This ensures that the resonator surface is exactly centred on the rotational axis and polishing can be accomplished by rubbing the polishing pad side-to-side while the resonator rotates beneath. Beginning with a relatively hard fabric polishing pad, such as optic cleaning tissue, and a slurry with a medium particle size, $\approx 0.5\mu$m, large chips or defects in the resonator can be removed if any were present.

For the polishing of brittle optical materials we use slurries containing polycrystalline diamond particles ranging in size from 0.05 μm to 3 μm. The slurries are in the form of a thick paste that can be applied to a polishing pad and are purchased from Allied High-Tech. We used folded optical cleaning tissue to provide a rough polish and switch to a dedicated synthetic polishing pad (Chem-Pol from Allied High-Tech) to provide a finishing polish.

A finishing polish is done by moving to smaller polishing particle sizes, down to ≈ 0.05 μm, then to medium particle sizes, ≈ 0.2 μm, with a softer polishing pad. We saw little improvement in the optical Q-factor by moving to particle sizes smaller than 0.1 μm on soft polishing pads.

The steps taken to produce a smooth, polished surface depend on the material being polished as well as the initial surface quality from the turning process. $\mathrm{CaF_2}$ whispering gallery resonators, for example, can be rapidly polished due to the soft material and the simple geometry of the resonators. In this case significant defects can be polished away with a rough polish and each subsequent step can take as little as a few minutes to remove the roughness of the previous polish. A square $\mathrm{LiNbO_3}$ resonator, on the other hand, requires a hand polish of each mirror surface and the removal rates are much slower. A good initial surface quality is therefore required and polishing can still take more than of a day.

During the polishing process, the surface finish of the resonator can be monitored using the on-machine white-light interferometer (see section 14.3). The polishing process starts with a step that is sufficient to remove the existing damage. For a ductile cut in $\mathrm{LiNbO_3}$ this is usually a 0.5 μm diamond size with a finishing polishing pad. A severely chipped $\mathrm{CaF_2}$ resonator could require an initial polish that uses 1 μm grit with a rough cleaning tissue. Once the surface roughness that is visible on the interferometer appears to be limited by the current polishing step, the grit size is reduced and the process is repeated. A minimum grit size of 0.1 μm seems to be sufficient to produce resonators that are not limited by roughness. Starting from a very rough polish, the step sizes used are: $1\mu\mathrm{m} \rightarrow 0.5\mu\mathrm{m} \rightarrow 0.25\mu\mathrm{m} \rightarrow 0.1\mu\mathrm{m}$ with optical tissue then $0.5\mu\mathrm{m} \rightarrow 0.25\mu\mathrm{m} \rightarrow 0.1\mu\mathrm{m} \rightarrow 0.05\mu\mathrm{m}$ with the synthetic polishing pad.

Figure 13.9 shows the effects of polishing on diamond turned resonators. The top row shows white-light interferometer images of a $\mathrm{LiNbO_3}$ resonator before (a) and after (b) polishing. The resonator was polished in a series of steps that used decreasing sizes of diamond particles, down to 0.05μm applied to a final polishing pad. The lower row shows similar images of a $\mathrm{CaF_2}$ resonator after a rough polish with $0.5\mu m$ diamond particles on an optical cleaning tissue. This shows that the tool lines have been removed but scratches from the rough polishing remain and need to be removed by further polishing.

The polishing pad or tissue carrying the diamond slurry was rubbed against the resonators by hand. A moderate pressure is exerted while moving the polishing pad and care needs to be taken to slightly randomise the pattern of motion. This can be taxing if the polishing takes more than a few hours. The task is greatly simplified for whispering gallery resonators. These were rotated by the lathe at $\approx$2000 rpm while the polishing pad was moved side-to-side, perpendicular to the direction of rotation. In this manner, material could be removed much more quickly than if done entirely by hand.

13.7 Non-cylindrically symmetric resonators

The total internal reflection resonators were fabricated similarly to the whispering gallery resonators, however the cutting depth was made to be a function of the spindle angle. This

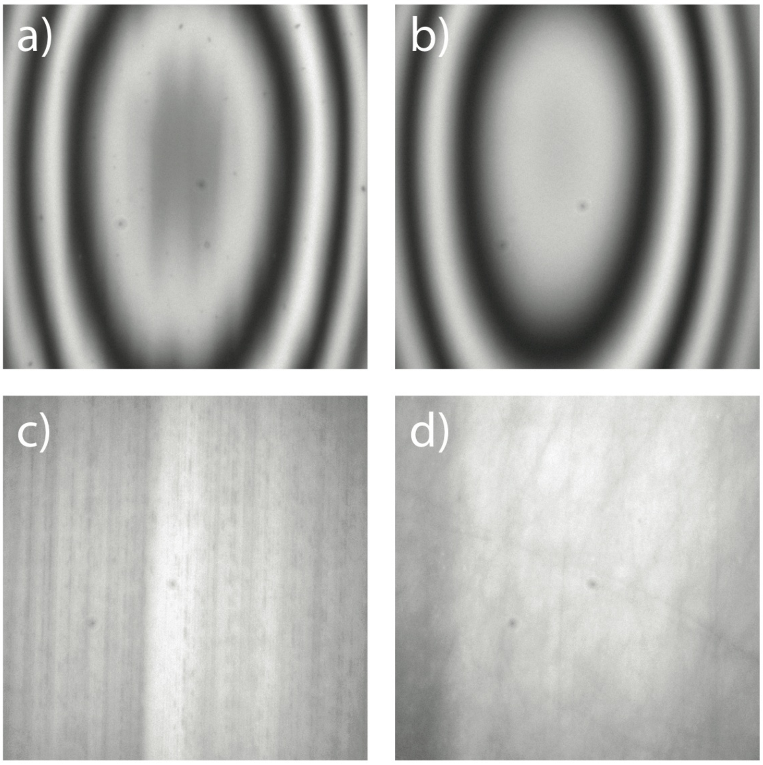

Figure 13.9: White-light interferometer images taken before (left column) and after (right column) polishing. The top row is a $LiNbO_3$ monolithic resonator and the bottom is a CaF_2 resonator.

necessitates slowing the rotation rate during cutting from ≈ 2000 revolutions per minute (rpm) to ≈ 60 rpm resulting in a much longer process. This type of cutting places results in a careful balance of parameters required to produce a resonator with acceptable surface finish and form accuracy.

Fabrication starts with mounting a blank in the same way as for producing whispering gallery resonators. The diameter of the blank is then reduced to a cylinder that is slightly larger than the final dimension of the resonator, while attempting to avoid causing significant chipping in the final few passes. The exact diameter of the blank is measured using an on-machine white-light interferometer so that this size can be used for programming the cut.

The non-cylindrical cut consists of roughing passes, semi-rough passes and finally finishing passes. Because the resonators need to have exact curvatures on each mirror face and because the diameter of the resonator needs to be known exactly to program the cuts, it is critical that the semi-rough and finishing passes cause no chipping so that no additional cuts are needed and only light polishing is required.

The three-fold symmetry of the fracture pattern in $LiNbO_3$, section 13.4, makes it difficult to avoid cutting at least one of the mirror faces in a region that is prone to brittle fractures. The only way to ensure that a ductile-mode cut occurs even in the fracture-prone regions is to use a nearly new diamond tool for the roughing passes and then switch to a fresh cutting edge, either an unused portion of the tool or a completely new tool, for the semi-rough and finishing passes.

The cut depths and number of cuts are a balance between removing enough material to form the mirror surfaces and to get below chipping from the roughing passes while not making so many cuts that the new tool becomes dull.

CHAPTER 14

Experimental Techniques

14.1 Lasers

Three lasers were used for the experiments with monolithic resonators. The primary laser used was a New Focus TLB-6721 Velocity tunable diode laser with an output power of 40 mW. This laser uses a Littman-Metcalf cavity configuration to provide a 22 nm continuously tunable mode-hop free range from 1050 nm to 1072 nm. The large tuning range of the laser and the capability of rapidly scanning over a 22 GHz fine-tuning range makes the Velocity laser ideally suited to experiments on small resonators featuring large free spectral ranges. Furthermore, a fast modulation input means that the laser can be directly modulated for locking purposes. The Ti:Sapphire laser that is discussed in section 7.2.1 was also used as a tunable source for some of our early investigations with CaF_2 whispering gallery resonators.

The other laser that was widely used was an InnoLight Diablo laser. This features a monolithic Nd:YAG resonator as the laser cavity for high-stability operation. The cavity design is a non-planar ring oscillator that uses the Faraday effect to break degeneracy for counter-propagating modes to achieve single-mode operation. The laser outputs 2 W of 1064 nm light and has an integrated frequency doubler. With the doubler operating, 100 mW of 532 nm light is available along with 0.8 W of 1064 nm light. Because of the monolithic lasing cavity the tuning rate and range of the laser are severely limited. The availability of the second harmonic and the high output power, however, make the Diablo necessary for investigating non-linear optical effects.

The two 1064 nm lasers are prepared into collimated beams of roughly the same size and are combined into a single spatial mode on a polarising beam-splitter. An anamorphic prism pair is used to circularize the elliptical output of the Velocity laser to improve the mode overlap between the two. The combined beams are then sent through a polarization maintaining fibre to ensure that they are in identical, Gaussian spatial modes. Placing the two lasers into a common spatial modes allows us to use either laser to probe the resonators without aligning each independently. A wave-plate is used to rotate the polarisation after the combining beam-splitter depending on which of the lasers is used so that a consistent polarisation is sent into the optical fibre.

Having the two lasers aligned to a single-mode fibre is useful because it enables us to perform alignments with the rapidly tunable laser prior to switching to the higher-powered Diablo laser for non-linear optical conversion. A sketch of the optical layout used

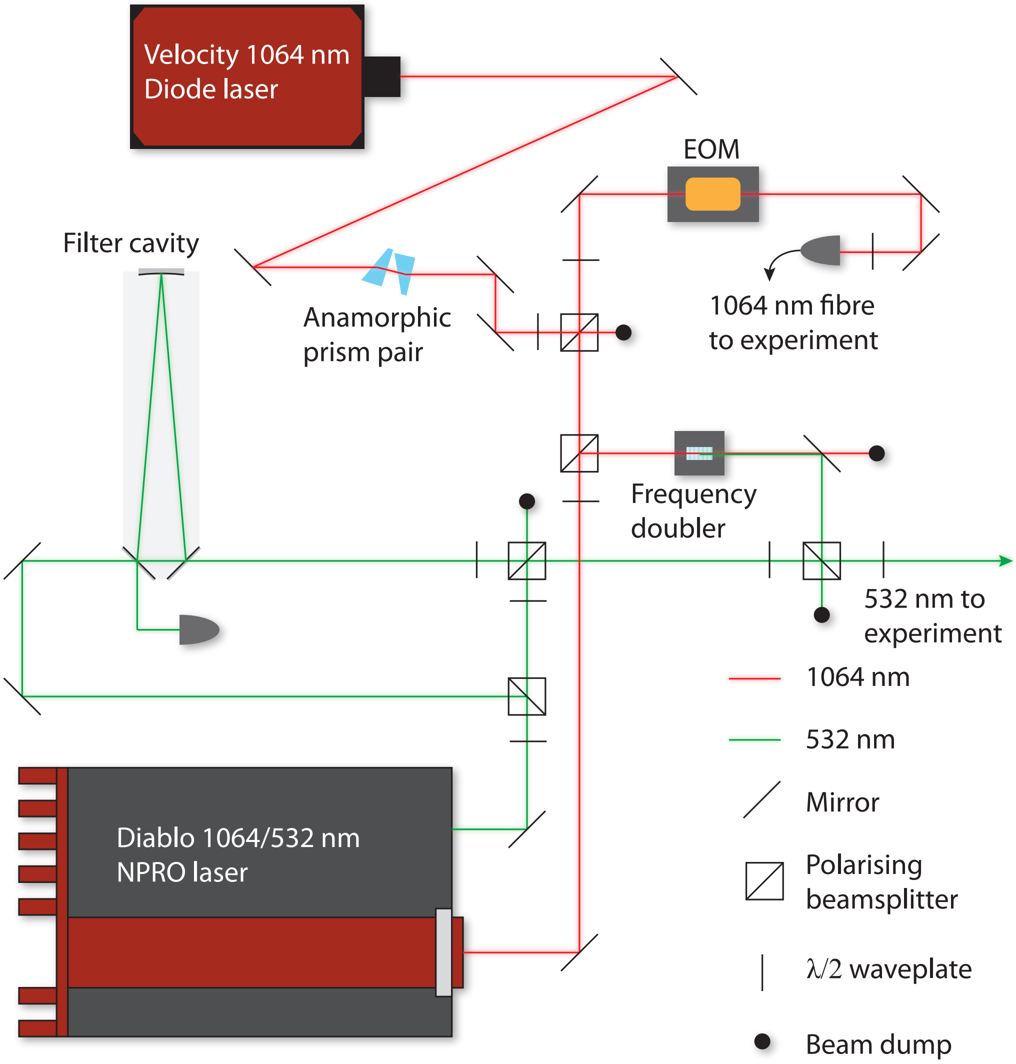

Figure 14.1: Layout of the optics for producing 1064 nm and 532 nm probes for experiments on monolithic resonators. The outputs from a tunable diode laser and an Nd:YAG laser are combined into a common spatial mode, as are two sources of second harmonic. The combined modes allow us to use the most frequency-agile source for alignment and switch to a higher power source for non-linear frequency conversion.

to combine the laser outputs is shown in fig. 14.1.

Similarly, multiple sources of second-harmonic light were combined into a single beam that could be used to measure the resonator properties at 532 nm or to drive a non-linear down-conversion process. The main source was the integrated doubler in the Diablo laser which provides roughly 0.8 W of 532 nm light. This is passed into a filter cavity to obtain a TEM00 mode. If the laser needs to be scanned over the resonator spectrum, the mode-cleaner is bypassed due to its limited free-spectral range. The integrated doubler will follow the fundamental through large frequency changes, however, the power level is not constant as the auto-lock feature adjusts the doubler to maintain or re-acquire a frequency lock with the fundamental. For a more stable, tunable source of 532 nm light, we used a single-pass optical frequency doubler, the output of which was also combined on a polarising beam-splitter with the other 532 nm beam. Only a few milliwatts were produced from this source but it was sufficient for linear measurements.

14.2 Mechanics

One of the main difficulties in working with monolithic resonators is the mechanical setup associated with positioning the coupling prisms near the resonator with the orientation required for evanescent coupling. Ideally, the prisms should be mounted such that the prism surface used for coupling is parallel to the surface of the resonator at the location that the optical mode is incident on the resonator surface.

The first design that was used was to have the resonator held at a fixed position, usually mounted on an oven, with the prism mounted on a translation stage with piezo fine-adjustments. The workhorse translation stage was the nanoMax 311 D/M three-axis flexure stage from Thorlabs. The flexure stage provided a platform that could be positioned with very low-backlash for repeatability. The stage is equipped with closed-loop piezo drives as well as differential micrometer screws for coarse positioning. Two of the axes were not strictly necessary, but were useful to move damaged or dirty areas of the prism away from the coupling region.

A mirror mount was placed on top of the three-axis stage to act as an additional two-axis rotation stage. This was used to align the coupling face of the prism to be parallel to that of the monolithic resonator. The disadvantage to the stage is that the physical dimensions of the stage, along with the size of the resonator oven, necessitated mounting the prism at the end of a cantilever that was fixed to the stage. The resulting combination of stages was flexible and had sufficient degrees of freedom but suffered from poor mechanical stability.

14.2.1 Mechanics for whispering gallery resonators

In the case of a whispering gallery resonator, the situation is simplified by the rotational symmetry of the resonator. Ensuring that the prism lies flat in the plane of the resonator is sufficient for good alignment. The only remaining task is then to carefully tune the distance between the surfaces involved in the evanescent coupling. We initially used the mechanical setup described above due to its flexibility. This allowed us to investigate the effect of the prism angle on the coupled mode shape. To improve the mechanical stability, however, we moved to a more compact design.

A simplified coupling setup that consisted of a single attoCube closed-loop piezo translation stage was designed. The stage provided both long travel and fine positioning in a single package. For this coupling setup, the resonator was machined on a mounting stud that was designed to screw into a baseplate to which the translation stage was attached. When fully screwed down, the stud sat flush with the baseplate and ensured that the plane of the resonator was aligned with the plane of the stage. A prism mounted on the stage can be brought into proximity of the resonator, simplifying the coupling process to a single alignment. The setup was effective in allowing resonators to be tested rapidly, however, it was found that the positioning of the attoCube was too coarse to be effective. This could be corrected by using a different controller for the stage. The controller that we had available has the capability to move the stage in discrete steps but not to fine-tune the position using a constant voltage to the piezo in the stage.

14.2.2 Mechanics for total internal reflection resonators

The design for the mechanics eventually used to couple to the total internal reflection resonators was an attempt to provide all the required degrees of freedom for coupling in a compact and stable package. The resonator was held on a central, fixed post with an integrated Peltier element for temperature control. Prisms could be attached at either side for input and output coupling or independent coupling of the fundamental and second harmonic. The prisms were mounted on custom designed piezo-actuated flexure stages which were, in turn, fixed directly to compact high stability mirror mounts from Radiant Dyes. These mirrors mounts were then attached to roller-bearing translation stages actuated by differential micrometer screws.

The roller bearing provided a long travel so that prisms could be easily removed or added to the setup away from the resonator and allow rapid coarse positioning. It was found that the added stiffness of roller bearing stages compared to flexure stages, along the with long travel was an advantage over the NanoMax design for stability and ease of use. The custom piezo actuated flexure provided fast fine positing in a compact area while the mirror mount provided the necessary two rotational degrees of freedom for aligning the prism relative to a flat resonator surface. The translation stages were finally mounted to a base plate that also supported the post for the resonator so the entire coupling setup could be moved as a single piece. The design made it straightforward to change coupling prisms by simply swapping out the entire translation stage with aligned optics in place.

The final element in the coupling setup was an adapter to which the mounting studs attached to. This adapter used a two-dimensional, square tongue-and-groove mating system to fit onto the fixed post in the experiment and onto a chuck adapter that is used during the resonator fabrication process. This provided a precise angle reference so that non circularly symmetric resonators could be exactly located relative to the angle of the coupling interface. This ensured that the coupling interface was nearly aligned with the prism face when initially placed in the experiment to simplify coupling.

14.3 White-light interferometer

Precision fabrication requires precision metrology to detect errors in produced parts. One key piece of equipment is the white-light interferometer. The microscope uses a beamsplitter in the objective to form an interference fringe between the part being measured and a reference flat. Illuminating the sample with low-coherence light using another beamsplitter in the microscope body produces an interference fringe only when the path length difference between the sample surface and the reference flat is on the order of a wavelength.

This configuration is known as a Mirau interferometer and is illustrated in fig. 14.2 and pictured in fig. 14.3. Scanning the distance between the objective and the sample and recording when the maximum bright fringe occurs for each point in the microscope image provides a height map of the surface that is being imaged.

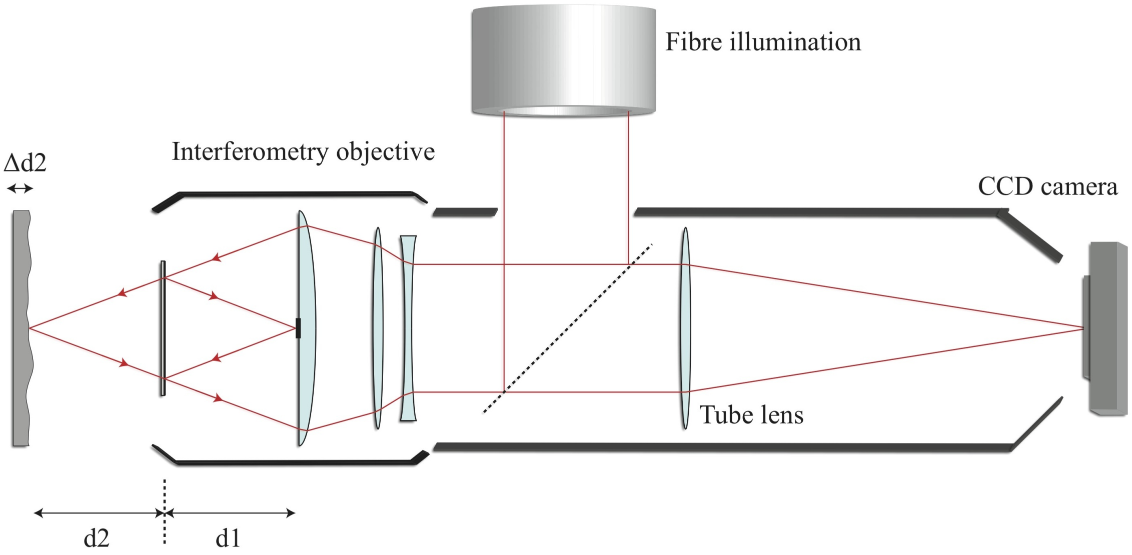

Figure 14.2: Schematic of the white-light interferometer. The Mirau interferometry objective is illuminated with broadband light. An interference fringe is visible only when the distances d_1 and d_2 are equal, providing a sub-wavelength measurement of the relative heights of each point on the sample.

We constructed a white-light interferometer directly on the lathe using off-the-shelf microscopy components. The objective used was a 20X Nikon CF IC Epi Plan DI interferometry objective. This was then mounted on an Infinitube FM-200 microscope body with an in-line fibre illuminator port. The white light sources used were either a halogen source for broadband light or a three-LED illuminator for narrow-band light at 465 nm, 521 nm or 635 nm. The distance between the objective and the work-piece was adjusted using the lathe's motion control with nanometre accuracy. A height map of the work-piece surface could be constructed by recording uncompressed, fixed-frame rate video of the fringes while scanning the objective distance at a known rate. The video was then imported to Mathematica and processed into a height map by finding the time at which the maximum fringe occurred for each pixel.

14.4 Temperature control

Temperature control for the resonators was done using a Wavelength Electronics LFI-3751 temperature controller. This controller stabilises the temperature of a monitored system using a PID control loop. The temperature sensor is a small glass-encased thermistor and temperature feedback is provided by a Peltier thermo-electric element located under the resonator. A number of designs were used for the mechanical construction of the resonator mount so the thermal design varied to some degree with the different designs. The general structure however, was that the cold side of the Peltier element was in thermal contact with an aluminium mounting structure that was directly attached to the optical bench.

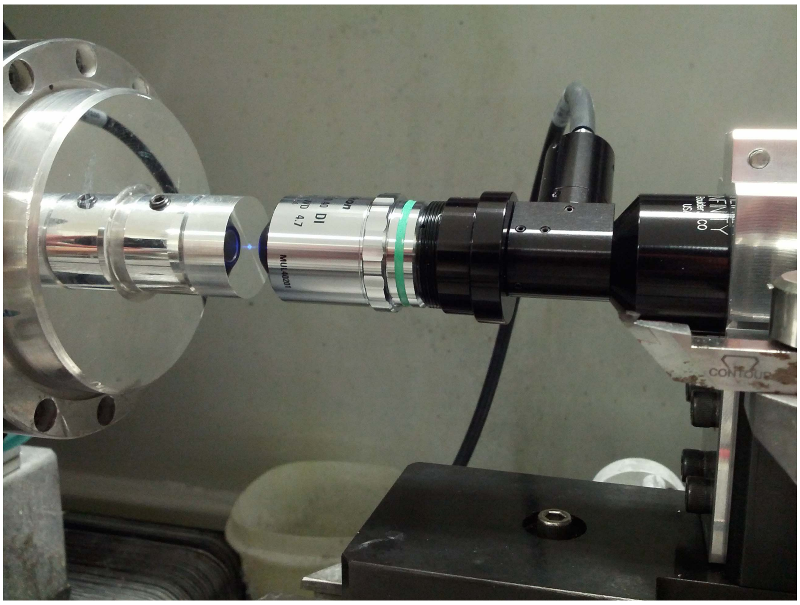

Figure 14.3: The interferometer mounted on the lathe.

Higher temperatures could be attained by insulating the cold side of the Peltier from the optical bench and heating the cold side of the element using a separate resistor, however, this was only necessary to heat the Peltier above 90° C, which was not required for 7% MgO doped $\mathrm{LiNbO_3}$ due to a lower phase-matching temperature. One design also used a three-axis rotation stage between the Peltier and the optical bench, however heating of the stage resulted in significant instability in the position of the resonator, which needed to be kept within evanescent coupling distance of the prism.

14.5 Evanescent coupling

The monolithic resonators are coupled to free-space optical modes through frustrated total internal reflection. This technique involves bringing a prism into the evanescent field of the resonator. Some of the field will then be transmitted into the prism and therefore to a free-space mode. The considerations for efficient coupling are spatial overlap between the out-coupled field and a field input to the coupling prism, the amplitude of the field that is transmitted through the evanescent interface, and the impedance-matching condition of the resonator. It is worth discussing each of these separately and we begin with the spatial overlap.

14.5.1 Spatial mode-matching

Following the method in [a39] we compute the field amplitudes for a resonator mode $A_0(t)$ and for a set of modes $B_k(t)$ within the coupler. The driving field which is incident on the coupler has a time-varying intensity $|B^{in}(t)|^2 = \sum_k |B_k^{in}(t)|^2$ distributed across N modes, each with an amplitude $B_k(t)$.

Each coupler mode is coupled to the resonator with a transmittance T_k. We define the vectors $\mathbf{B}$ and $\mathbf{T}$ to denote the amplitude and transmittance in each coupler mode, respectively. The internal reflectivity, R, is given by the coupling of the resonator to all coupler modes and, in the limit of $R \ll 1$, is

$$\begin{aligned} R &= \prod R_k = \prod \sqrt{1-T_k^2} \\ &\simeq \prod (1-T_k^2/2) \simeq 1 - \sum T_k^2/2 \\ R &\simeq 1 - T^2/2. \end{aligned} \tag{14.1}$$

The amplitudes of the resonator mode and the modes at the output of the coupler are given by

$$\begin{aligned} A_0(t) &= i\sum T_k B_k^{in}(t) + RA_0(t-\tau_0)exp(i2\pi n_r L/\lambda - \alpha L/2) \\ B_k^{out}(t) &= R_k B_k^{in}(t) + iT_k A_0(t), \end{aligned} \tag{14.2a}$$

where $L \simeq 2\pi a$ is the path length of the resonator mode for a resonator of radius a, n_r is the refractive index of the resonator, λ is the optical wavelength, α is the linear attenuation and $\tau_0 = n_r L/c$ is the circulation time of light in the resonator. If the losses per round-trip are small, we can assume that the amplitude of the resonator mode varies slowly compared to τ_0 and write the evolution as

$$\frac{dA_0}{dt} = iCB^{in} - (\delta_c + \delta_0 + i\Delta\omega)A_0, \tag{14.3}$$

where

$$\delta_0 = \frac{\alpha c}{2n_r}; \delta c = \frac{T^2}{2\tau_0}; C = \frac{T\Gamma}{\tau_0}. \tag{14.4}$$

The coefficient

$$\Gamma = \frac{\mathbf{T}\cdot\mathbf{B}^{in}}{TB^{in}} \leq 1 \tag{14.5}$$

describes the mode-matching between the incident light and out-coupled light. δ_0 is the intrinsic loss rate and δ_c is the loss rate through coupling with the propagating modes. The steady-state solution to eq. (14.3) is

$$A_0 = \frac{iT\Gamma B^{in}}{\tau_0(\delta_c + \delta_0 + i\Delta\omega)} = \frac{i2\delta_c B^{in}}{\delta_c + \delta_0 + i\Delta\omega}\frac{\Gamma}{T} \tag{14.6}$$

and the resultant field at the output of the coupler is

$$\mathbf{B}^{out} = \mathbf{B}^{in} - \frac{B^{in}\Gamma}{T}\frac{2\delta_c}{\delta_0 + \delta_c + i\Delta\omega}\mathbf{T}. \tag{14.7}$$

Equation (14.7) provides a framework to discuss how coupling to an optical resonator will affect a general monochromatic input optical field. We can choose an orthogonal set of optical modes to which the resonator mode couples. In a standard, non-astigmatic, optical resonator consisting of mirrors separated by free space the modes of the resonator match the Hermite-Gauss modes that are typically produced by lasers. This is the case commonly encountered while mode-matching to cavities in the lab and 14.7 can be reduced to a single-mode coupling between a TEM cavity mode and a free space mode with matched waist and divergence. Another simple case is fibre-coupled cavities in which only one coupler mode exists.

In general, however, if the mode of the input laser isn't matched to that of the resonator, a spatial redistribution of optical power will occur as a result of coupling to the resonator. Here, the coupling coefficients between the resonator mode and each of the modes in the coupler needs to be worked out. While this is not necessarily straightforward to compute, it does provide a theoretical description of how the mode shape of an input beam changes as it interacts with the resonator.

14.5.2 Evanescent transmission

Prism coupling to a monolithic resonator works on the principle of frustrated total internal reflection (FTIR). Total internal reflection occurs when light, propagating from a material with a higher index of refraction, n_1 to one of a lower index, n_2, is incident on the boundary at an angle larger than the critical angle. In this case, Snell's law,

$$n_1 \sin\theta_1 = n_2 \sin\theta_2, \tag{14.8}$$

cannot be satisfied with a real angle, θ_2, and the light is entirely reflected from the boundary. The electromagnetic field, however, must be continuous across the boundary, which is satisfied by an exponentially decaying field on the lower index side of the boundary, called the evanescent field. The magnitude of this evanescent field [b40] is

$$\begin{align} E(z) &= E(0)e^{-\beta z} \tag{14.9} \\ \beta &= \frac{2\pi}{\lambda_0}\sqrt{(n_1 \sin\theta_1)^2 - n_2^2}. \tag{14.10} \end{align}$$

Bringing a second dielectric into the evanescent field will result in an electric field in that material and, if the refractive index is high enough, it will form a propagating field. This can be used to couple light across a thin region of lower index material, analogously to quantum tunnelling. The Fresnel equations [b40],

$$\begin{align} r_p &= \frac{n_2\cos\theta_1 - n_1\cos\theta_2}{n_2\cos\theta_1 + n_1\cos\theta_2} \tag{14.11a} \\ t_p &= \frac{2n_1\cos\theta_1}{n_2\cos\theta_1 + n_1\cos\theta_2} \tag{14.11b} \\ r_s &= \frac{n_1\cos\theta_1 - n_2\cos\theta_2}{n_1\cos\theta_1 + n_2\cos\theta_2} \tag{14.11c} \\ t_s &= \frac{2n_1\cos\theta_1}{n_1\cos\theta_1 + n_2\cos\theta_2}, \tag{14.11d} \end{align}$$

give the reflected and transmitted amplitudes for an electromagnetic wave incident on a boundary between two dielectrics. From these, we can calculate the reflection and transmission in the case of frustrated total internal reflection.

Following [a41], we consider the amplitude of the light reflected from the first boundary as it is evanescently coupled from a layer of material with refractive index n_1 to one with index n_3 through an intermediate layer with index n_2. As the light is initially incident on the first boundary, a portion of the amplitude, r_{12}, is reflected. Some of the transmitted light is further reflected from the $2 \to 3$ boundary and is subsequently transmitted back through the $2 \to 1$ boundary with an amplitude $t_{12}e^{i\delta}r_{23}t_{21}$, where $\delta = (4\pi d/\lambda)\sqrt{n_2^2 - (n_1\sin\theta)^2}$ is the phase accumulated due to the path length taken between the interfaces. Continuing with this reasoning, the total amplitude reflected from the 1 → 2 boundary is

$$r_{eff} = r_{12} + t_{12}t_{21}\left(\sum_{k=0}^{\infty} r_{21}^{k} r_{23}^{k+1} e^{i(k+1)\delta}\right). \tag{14.12}$$

Making use of the Stokes' relations,

$$t_{12}t_{21} + r_{12}^2 = 1 \tag{14.13}$$

$$r_{12} = -r_{21}, \tag{14.14}$$

and summing the geometric series, we obtain

$$r_{eff} = \frac{r_{12} - r_{32}e^{i\delta}}{1 - r_{12}r_{32}e^{i\delta}}. \tag{14.15}$$

The result is expressed in terms of r_{32} rather than r_{23} because we can easily calculate this coefficient. r_{32} corresponds to light propagating backward along the path, entering from n_3 and passing into n_2, and the output angle is easily computed: $\theta_{out} = \arcsin\left(\frac{n_1}{n_3}\sin\theta\right)$. Both r_{12} and r_{32} will be associated with a total internal reflection so we can rewrite them as $r_{12} = e^{i\gamma_1}$ and $r_{13} = e^{i\gamma_3}$ and, from eq. (14.10), define $\beta = i\delta/(2d)$. We calculate $T = 1 - |r_{eff}|^2$ and find

$$T(d) = \frac{2\sin(\gamma_1)\sin(\gamma_3)}{\cosh(2\beta d) - \cos(\gamma_1 + \gamma_3)}. \tag{14.16}$$

Figure 14.4 shows the intensity transmission, $T(d)$, through a small gap of lower index for various input angles. At low angles, Newton's rings are visible, while at higher angles evanescent coupling occurs. For comparison, eq. (14.10) scaled by the Fresnel transmittance for 45° is shown as the blue trace. The full calculation reveals that the decay follows a $(\cosh d)^{-1}$ dependence [a5, a41, a42] as opposed to a e^{-d} dependence for the simpler approximation.

Evanescent coupling results in a spatially non-uniform reflectivity where light is transmitted from the resonator to the coupler. This has two consequences: the spatial mode of the resonator is perturbed [b43] and the mode-shape within the resonator is not the same as the mode-shape on the surface of the coupler.

14.5.3 Critical coupling

If the input beam can be mode-matched to the out-coupled light from a resonator, eq. (14.7) reduces to

$$t = \left(1 - \frac{2\delta_c}{\delta_0 + \delta_c + i\Delta\omega}\right), \tag{14.17}$$

where, t is the amplitude transmittance. On resonance, the light is fully absorbed into the resonator if $\delta_c = \delta_0$, in which case the coupling loss is impedance-matched to the intrinsic loss. This point is referred to as critical coupling.

In the case where the input light is not mode-matched to a resonator mode, it is not possible to couple all incident light into the resonator. Figure 14.5 shows the intensity transmission as a function of detuning and coupling rate in the case of ideal mode-matching (black trace) and 50% overlap between the input light and out-coupled light.

An important point for experimentally observing critical coupling is that all of the

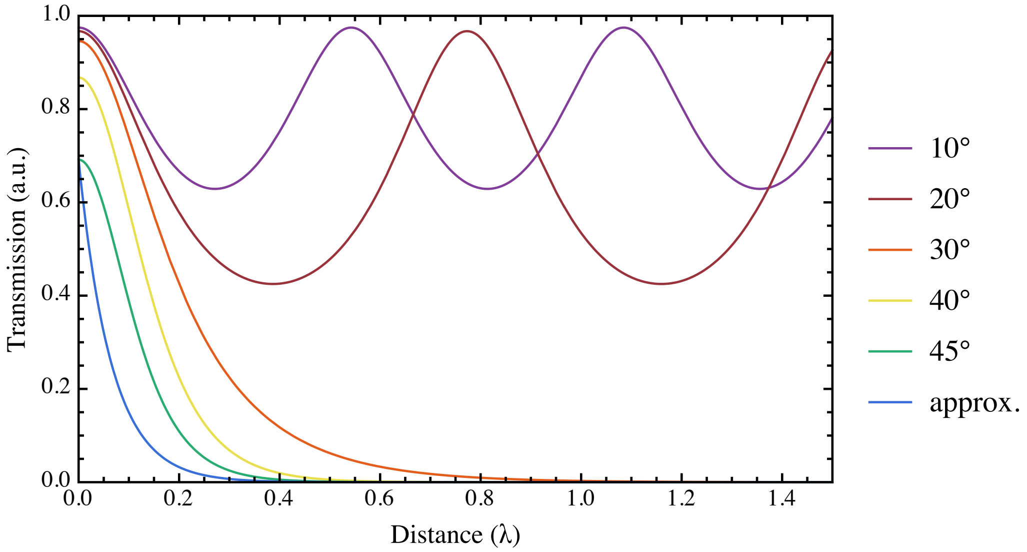

Figure 14.4: Intensity transmission between high index materials with a small gap between them. The indices of refraction are 2.23, 1, and 1.64, corresponding to coupling light from a monolithic resonator as described in chapter 16. The blue trace shows an approximation that considers only the decay of the evanescent field.

output light must be collected into the detector. In the case where the input light is not mode-matched to the resonator output, the spatial redistribution of optical power caused by the resonator will result in a lower detection efficiency near resonance and may appear as a higher coupling efficiency than is actually present. When the resonator is over-coupled, power is distributed into other free-space modes and can lead to dark fringes in the output mode even if the overall mode-overlap is quite poor. If the detector is located in one of these fringes, the coupling efficiency will appear to be artificially high.

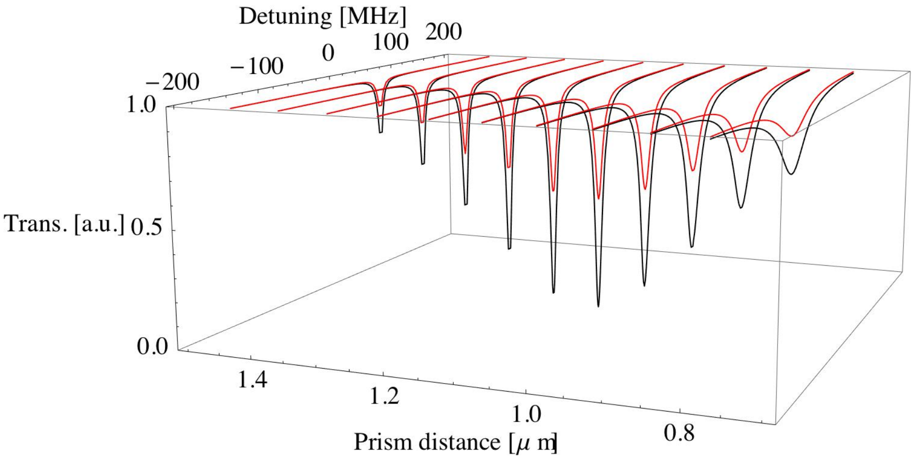

Figure 14.5: Plot of eq. (14.7) for various coupling rates. The back line corresponds to perfect mode-matching between the input and output modes, reducing to eq. (14.17) and the red line shows the case for mode-matching with 50% efficiency.

CHAPTER 15

Whispering gallery resonators

Whispering gallery resonators are a promising platform for producing compact optical devices, largely because the cylindrical symmetry makes it possible of fabricate very small, high finesse cavities. We produced a variety of such resonators from CaF_2 , MgF_2 and $LiNbO_3$. The resonators fabricated from CaF_2 and MgF_2 can reach high finesses (nearly 5000). Because the monolithic nature of the resonators makes them insensitive to vibration, they have applications as stable frequency references. The $LiNbO_3$ whispering gallery resonators are also tunable by the application of an external voltage, with applications in optical filters and high-frequency reference oscillators [b1]. We have also made preliminary progress towards non-linear conversion in whispering gallery resonators, although the resonators that will be presented in chapter 16 are more well developed for stable non-linear processes.

15.1 Alignment

The alignment procedure for whispering gallery resonators involves tuning the prism distance to achieve evanescent coupling, finding the angle for the input beam that phase-matches the evanescent field of the resonator with that on the prism base and shaping the input beam to match an output mode of the resonator.

The resonator is placed near the centre of the prism base with the smallest gap that is still easily resolvable by a video microscope. Figure 15.2 shows a microscope image of a resonator that is almost contacting the prism, although the gap is too small to be seen in the image. The vertical alignment of the input beam is done first, such that the beam lies in the plane of the resonator and at the same height as its apex. The focusing lens is adjusted to this height as well so that the vertical alignment is roughly fixed. The input beam is then sent into the coupling prism, without a focusing lens, at an angle that is smaller (relative to a line normal to the prism base) than the calculated phase-matching angle for the resonator. The beam is aligned to strike the prism base in the location of the resonator. This ensures that a large number of modes can be excited.

The focusing lens is then inserted into the beam such that it focuses the beam onto the prism base at the location of the resonator. The location of the totally internally reflected input beam on the prism base can be seen with a video microscope or infra-red viewing scope, or by eye for visible light. The resonator is gently advanced towards the prism using a translation stage. For $LiNbO_3$ resonators, contact with the prism had no effect on the

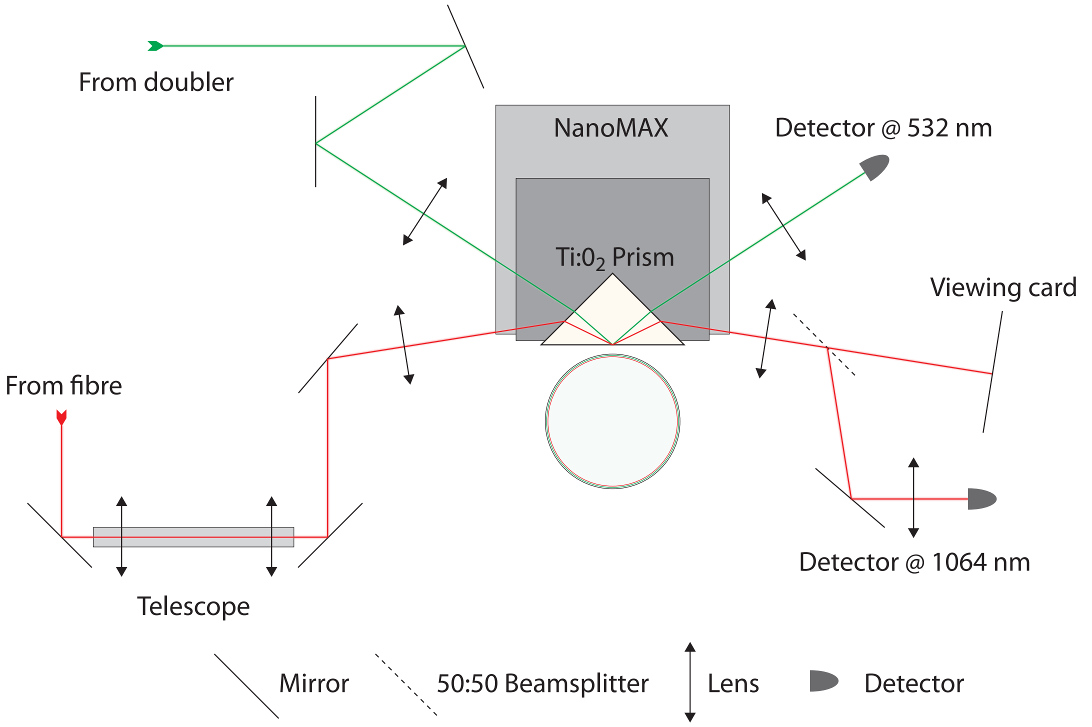

Figure 15.1: Layout of the optics for probing whispering gallery resonators.

finesse of the resonator if it was advanced until contact was made. CaF_2 resonators are more fragile (although they are also generally unaffected by a few gentle contacts with a prism) so the resonator is advanced in wavelength-size steps while performing the optical alignment repeatedly in hopes of obtaining some coupling before contact.

When the resonator is in contact with the prism or close enough that coupling may be visible, the shape of the output beam is monitored while adjusting the input alignment. Most of the initial alignment is done by freely moving the focusing lens. As the beam is nearly focused on the prism base, defects in the base will be imaged in the far field and can be seen on a viewing card. Each defect can be inspected by focusing the beam onto it by adjusting the position of the focusing lens. Scanning the laser frequency will result in a moving fringe pattern between the out-coupled resonator modes and the reflected input beam when the beam is focused and so the presence or absence of these fringes will inform whether a given defect has evanescent coupling to the resonator or to dust particles on the prism.

Once the location of the prism base has been identified in the image formed on the viewing card, the position of the lens can be adjusted to focus onto the coupling region. This is done by moving the lens so that the image of the prism base fills the beam. Concentric interference fringes that move as the excitation laser frequency is scanned indicate how well the divergence of the input beam is matched to that of the out-coupled light, just as they would for mode-matching two free-space beams.

Once the excitation beam is roughly aligned, the alignment can be fine-tuned by walking the input beam with the mirrors that are located before the coupling lens, monitoring both the absorption spectrum and the mode shape on the viewing card. A schematic

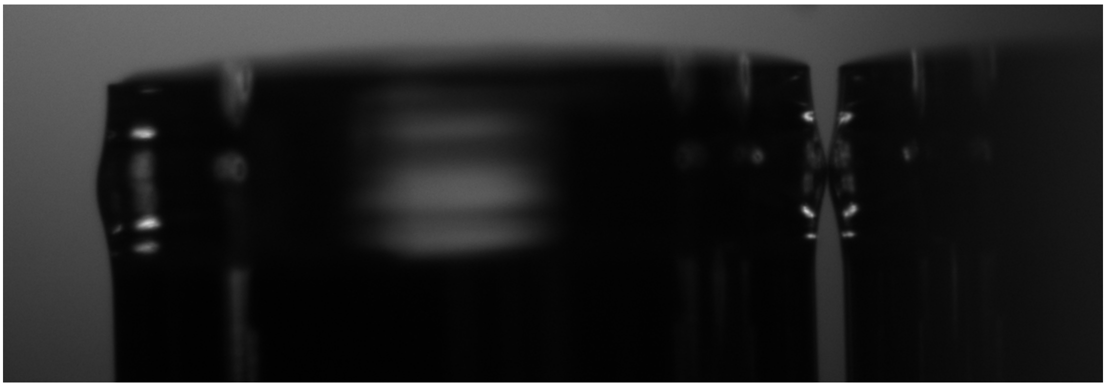

Figure 15.2: Prism coupling the smallest of the resonators that we produced, a $LiNbO_3$ whispering gallery resonator with a 1.4 mm diameter and 200 μm polar radius. The reflection of the resonator (on the right of the image) can be seen in the prism.

of the optical layout for coupling is shown in fig. 15.1. A telescope located before the mode-matching optics can be used to adjust the beam size to match that of the resonator modes. Astigmatic optics (cylindrical lenses and anamorphic prism pairs) were used prior to the telescope to attempt to match the astigmatism of the out-coupled light with limited success. Efficiencies only up to 80% were obtained in both $LiNbO_3$ and CaF_2 .

15.2 Prism coupling for whispering gallery resonators

To couple light into the whispering gallery resonators we used prisms with a higher refractive index than the resonator material. If the index of the coupling prism is less than that of the resonator, the prism will not frustrate the total internal reflection and no light will be coupled to free-space. Even with a higher index coupling prism, in the limit of tangential propagation, the coupling from the resonator to nearby dielectrics vanishes according to the Fresnel equations, (14.11).

The magnitude and spatial mode of the light coupled from a whispering gallery resonator can be calculated by considering the resonant mode itself, and from that, the evanescent field associated with it. This can be done for small, spherical resonators [a39], however the difficulty in computing the mode structure for larger, aspheric resonators makes it difficult to determine the coupling pattern.

Without attempting to model a given resonator completely, a best guess for achieving a high coupling efficiency is to match the shape of the evanescent coupling region to a free-space beam. If the size of a given mode is larger than the coupling region, the out-coupled shape will be be determined primarily by the shape of the coupling. In the absence of a coupler, the evanescent field will decay exponentially away from the resonator while the distance between the resonator and coupler will increase quadratically moving away from the point of closest approach. These dependencies combine to form a Gaussian coupling aperture.

The resonator can be cut such that the aperture results in a circular out-coupled beam by tailoring the ratio between the resonator radius and the radius of curvature in the transverse direction. When the resonator and prism are in contact, the gap between the resonator surface and prism, d, will be

$$d(x, y) \simeq \frac{1}{2R^2}\left(x^2 + y^2\right), \tag{15.1}$$

where R is the radius of the resonator, r is the radius of curvature in the polar direction and x and y are Cartesian coordinates in the plane of the prism base. The angle that the out-coupled light will take in the coupling prism is given by Snell's law, eq. (14.8). Light incident on the coupling surface from within the resonator is nearly tangent to the surface, so $\sin\theta_1 \simeq 1$ and

$$\theta_2 \simeq \arcsin\left(\frac{n_1}{n_2}\right). \tag{15.2}$$

The footprint of the coupling region on the prism base will then have a circular aspect ratio when the ratio of r/R, is

$$\begin{aligned} r/R &= (\cos\theta)^2 \\ &= 1 - \left(\frac{n_r}{n_p}\right)^2. \end{aligned} \tag{15.3}$$

This method has been used to obtain a coupling efficiency of 99.96% [a44] efficiencies of 97%-99% have also been reported [b1, a13].

For oblate resonators, the shape of the out-coupled mode is not fully determined by the aperture but the internal shape of the mode has a significant impact on the out-coupled shape. This can be seen in the interference pattern between the input and output light and also in the emission patterns if a second prism is used as an out-coupler. The alignment of the input beam therefore has an effect on the relative coupling efficiencies to the modes. For prolate resonators, the evanescent aperture is the dominant effect in determining the out-coupled mode. In this case, the internal mode structure primarily determines the angle of the output mode, however, numerous modes may have a significant spatial overlap.

15.3 Mode structure

The mode structure of whispering gallery resonators is challenging to calculate. Analytic methods can determine the structure for spherical resonators but do not apply to the more common situation of toroidal resonators. The structure can be found by numerical simulations, however, which reveal the character of the modes. A complication to comparing the calculated mode shapes to experiments is the fact that the evanescent coupling region will introduce a windowing effect on the out-coupled mode. For the oblate resonators used to provide a roughly circular output mode, the out-coupled light has a shape that is more reflective of the geometry of the coupling region than the shape of the whispering gallery mode.

A large non-linearity between two optical modes requires the modes to be in nearly the same spatial mode. Because the out-coupled mode shape provides few clues as to what the internal resonators mode shape is, it is difficult to select modes for the fundamental and second-harmonic that have a large non-linearity. Furthermore, the rich spectrum of the whispering gallery resonators makes searching through all possible modes to find the pair of modes for the second-harmonic and fundamental with the highest non-linear interaction tedious and challenging. It should be noted that, while difficult, identification of particular modes in such geometries is possible. It has been shown that careful measurement of both

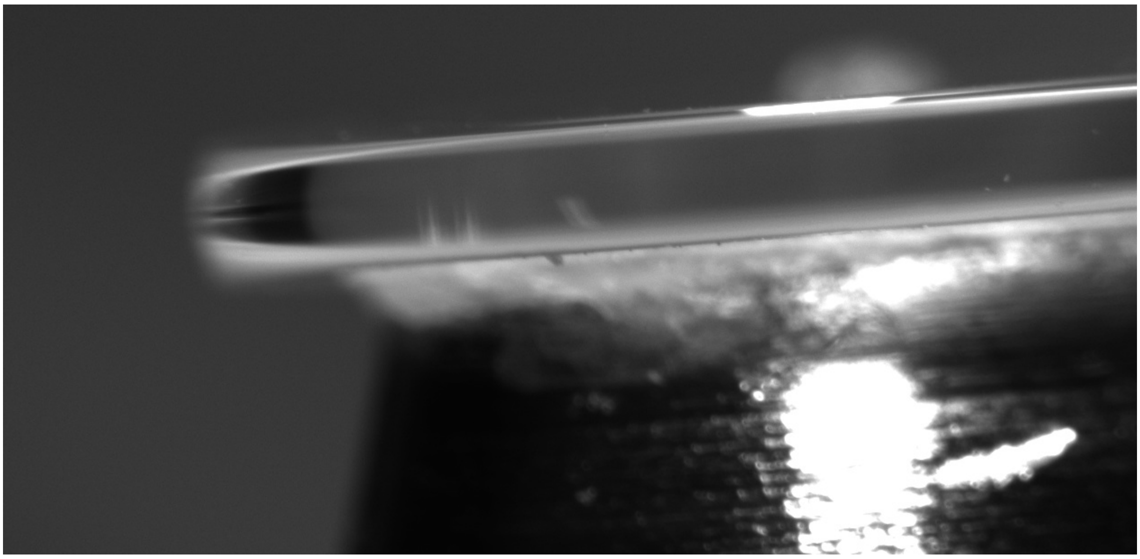

Figure 15.3: Optical microscope image of a calcium fluoride resonator.

the output patterns and spectrum can be used to infer the mode numbers associated with particular coupled modes [a45].

15.4 Results in CaF$_2$

The first material that we used for experiments on whispering gallery resonators was CaF$_2$. This was due to the ease of machining the material (see section 13.5) due to its softness compared to LiNbO$_3$ and due to very high optical Q-factors being reported in literature.

The CaF$_2$ blanks were diced from bulk rods, mounted on standard 1/2" optical posts using epoxy and machined with a ductile-mode finishing cut. Mounting them on standard posts enabled the resonators to be easily moved between the lathe and the optical measurement setup for testing. Various procedures for polishing were tested by directly measuring the unloaded Q-factor after each polish, resulting in the polishing schedule outlined in section 13.6. A microscope image of a finished resonator is shown in fig. 15.3.

Coupling was accomplished using a BK-7 glass prism and measurement of the resonator spectrum was done using a Coherent MBR series Ti:Sapphire laser operating at 795 nm. Figure 15.4 shows a typical spectrum obtained from a polished resonator. The maximum coupling efficiency obtained for this particular resonator was $\approx 70\%$ partly because it was not cut according to eq. (15.3).

The data shown in fig. 15.4 shows a characterisation of a $\approx$12 mm diameter CaF$_2$ whispering gallery resonator with a polar radius of curvature of 1 mm. Identifying modes that are repeated within the spectrum allows us to determine the free spectral range, 8.23 GHz, which is consistent with a 12.2 mm diameter resonator. Variation in the depth of the peaks is likely due to variations in the prism distance during the laser scan.

The linewidth of each resonance can be determined by fitting a Lorentzian, as is illustrated in fig. 15.4 (b) by the blue fit line to the data shown in red. This resonance has a loaded linewidth of 9.1 MHz, corresponding to a Q-factor of 4.2×10^7. Increasing the prism distance allows us to determine the intrinsic Q-factor by under-coupling the resonator. Fitting the under-coupled line, shown in (c), gives a linewidth of 1.7 MHz, corresponding to

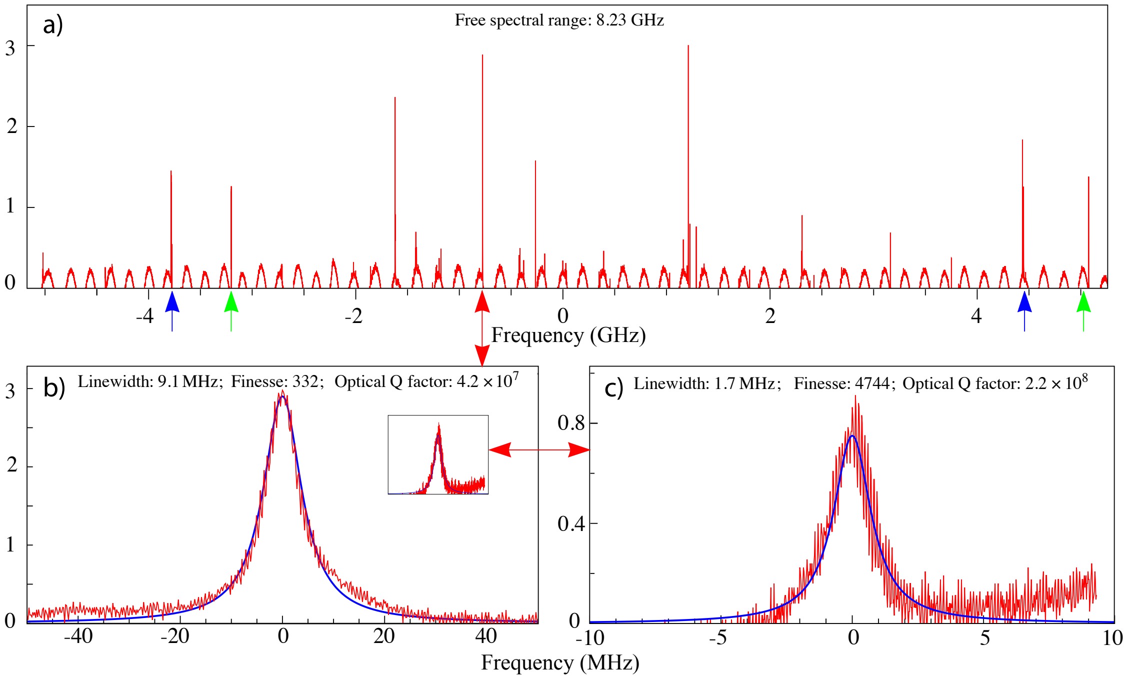

Figure 15.4: a) A spectrum obtained for a $\approx$ 12 mm diameter, CaF$_2$ radius whispering gallery resonator. The blue and green arrows identify repeated modes. b) Zoomed trace of the resonance in (a) highlighted by the red double-arrow. c) The same resonance as in (b), but undercoupled. This is also shown as an inset in (b), plotted on the same scale as the overcoupled line.

an unloaded Q-factor of 2.2×10^8 and a finesse of 4700.

15.5 Results in LiNbO$_3$

LiNbO$_3$ whispering gallery resonators, such as the one shown in fig. 13.5, were made with the goal of observing non-linear effects and, ultimately, engineering the resonators to produce continuous-wave vacuum squeezing. The two main technical challenges for the produced resonators was to obtain a high finesse and a high coupling efficiency, both required for efficient non-linear processes.

As with CaF$_2$, the circular crystal blanks were produced by dicing rods of bulk crystal. These were adhered to a brass mounting stud, as described in section 13.1, and thinned to a thickness of roughly 200 μm. A brass disk was then epoxied to the top to act as an electrode and was thinned to a few hundred μm. The assembly was machined to roughly the intended diameter using roughing cuts. As the diameter decreased, shallower cuts and sharp sections of the tool nose were used to limit the cutting forces on the blanks so that it didn't exceed the strength of the epoxy bond.

The resonators were cut as a bulge centred about the middle of the LiNbO$_3$ wafer (see fig. 13.5) with a radius of curvature defined by eq. (15.3). It was found that a ductile-mode cut could be produced in LiNbO$_3$ (see section 13.4), however, it was often simpler in practice to obtain a finishing pass with some small surface damage and remove that damage by polishing. This resulted in considerably less tool wear than attempting to obtain ideal surface finishes for each resonator.

The coupling setup, shown in fig. 15.1, is similar to the set-up for CaF$_2$. For exper-

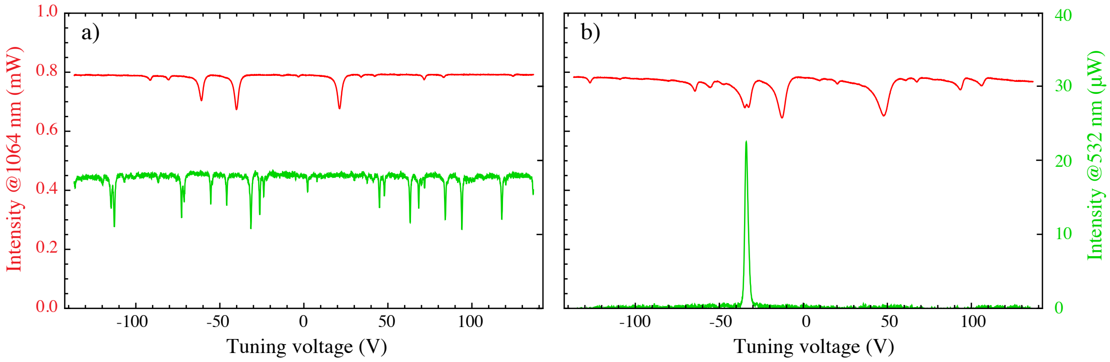

Figure 15.5: a) Spectra obtained for 1064 nm light (red trace) and 532 nm light (green trace) for a 2 mm diameter $LiNbO_3$ whispering gallery resonator. The resonator is critically coupled for the green light and over-coupled for the red light. Both wavelengths have a coupling efficiency of only about 40%. b) Second harmonic generation observed by blocking the green input and tuning the temperature until double-resonance is achieved.

iments with $LiNbO_3$ we used a NanoMAX three-axis positioning stage to position the prism next to the resonator. The axes in the plane of the prism base could be used to locate the contact point with the resonator in the centre of the prism base. The remaining axis adjusted the coupling distance.

Considerable effort was put into obtaining high coupling efficiencies; however, the best obtained were around 70%. The primary strategy was to adjust the ratio of the azimuthal and polar radii of curvature of the resonator to account for both the aperture effect of the evanescent coupling and the mode size of the whispering gallery mode itself.

The out-coupled mode shape was examined both by the interference with the input beam and by using a second prism to out-couple resonator modes and observe the mode shape directly. The aperture matching condition, eq. (15.3), produced no worse results than attempting to model and account for the mode size. This is consistent with the excellent results reported in literature, if somewhat surprising. The best coupling efficiencies were obtained slightly away from the critical angle, suggesting that a good mode-matching could be obtained by selecting a deeper mode to correct for an incorrect ellipticity. This seems to be consistent with the results reported in [b1].

In spite of best efforts to achieve a high coupling efficiency, visibilities above 70% were never obtained. Some non-linear effects were observed, however. Figure 15.5 shows the resonator spectra (a) of the fundamental (red) and second harmonic (green) as well as evidence of second harmonic generation (b). The spectrum of the resonator was scanned by applying a tuning voltage across the resonator while the laser frequency was held constant. In (a), both the fundamental and second harmonic are input to the resonator to observe the spectrum when the phase matching condition is not satisfied. Blocking the input of the second harmonic and tuning the temperature to obtain a phase-matched double resonance (b) demonstrates second harmonic generation. Some depletion of the fundamental can be seen where double-resonance occurs. In (a), the second harmonic is critically coupled while in (b) it is over-coupled so that generated light can exit the resonator. The fundamental is over-coupled in both cases.

CHAPTER 16

Total internal reflection resonators

Due to the difficulties associated with coupling efficiently to whispering gallery resonators, we decided to attempt an alternative geometry for the resonators: square resonators that support Hermite-Gauss modes. The resonators feature two flat sides for evanescent coupling so that the shape of the mode is preserved and two curved sides to create a stable cavity. The curved sides have a toroidal profile rather than a spherical profile so that the astigmatism due to a non-zero angle of incidence can be corrected. The curvature of these sides is chosen to maximise the non-linear interaction between a fundamental mode and a phase-matched second harmonic.

16.1 Designing the tool path

The shape of a square monolithic resonator consists of four toroidal mirror surfaces that are arranged in a square so that a square optical mode, rotated by 45° relative to the square of the resonator, is supported as illustrated in fig. 16.1(a). Each mirror has a different radius of curvature in the sagittal and tangential planes, allowing the astigmatism of the mode to be engineered.

It is convenient to discuss the resonator geometry in cylindrical coordinates, (r, θ, z). The tool path is programmed as a cut depth, d, relative to the apex of the resonator, R, as measured from the apex of one mirror to the apex of the opposing mirror. If the surface s describes the resonator surface, the cut depth will be $d(\theta, z) = s(\theta, z) - R$. We use a small angle approximation to consider the curvature in the sagittal and tangential planes independently, defining a cut depths

$$d_t(r_t, x_\theta) = \sqrt{r_t^2 - x_\theta^2} - r_t \tag{16.1}$$

$$d_s(r_s, z) = \sqrt{r_s^2 - z^2} - r_s. \tag{16.2}$$

The cut depth for each mirror is defined in Cartesian coordinates, (x_θ, z), rotated to be tangential to the apex of that mirror. The surface of the mirror is defined by the total of the two cut depths $d(x_\theta, z) = d_t(x_\theta) + d_s(z)$. A coordinate transformation to polar coordinates is performed on the tangential cut depth, with the correct offset in θ and r

for that segment. The surface is then defined by the piecewise function

$$d(\theta, z) = \left\{ \begin{array}{lcl} d_s(r_{s,0}, z) + T_0(d_t(r_{t,0}, x), \theta) & : & -\pi/4 < \theta \leq \pi/4 \\ d_s(r_{s,\pi/2}, z) + T_{\pi/2}(d_t(r_{t,\pi/2}, x), \theta) & : & \pi/4 < \theta \leq 3\pi/4 \\ d_s(r_{s,\pi}, z) + T_{\pi}(d_t(r_{t,\pi}, x), \theta) & : & 3\pi/4 < \theta \leq 5\pi/4 \\ d_s(r_{s,3\pi/2}, z) + T_{3\pi/2}(d_t(r_{t,3\pi/2}, x), \theta) & : & 5\pi/4 < \theta \leq 7\pi/4, \end{array} \right\} \tag{16.3}$$

where T_θ is a coordinate transformation for the tangential component for the mirror centred at θ.

Equation (16.3) describes a surface for the entire resonator; however, it is discontinuous unless each of the sides has identical parameters and, in any case, has a discontinuous derivative with respect to θ. To fix these problems, the tool path is further subdivided into an eight-part piecewise function, four surfaces, each occupying a $\pi/4$ slice centred about the apex of a mirror, and four interpolated regions. The interpolated regions are generated as the surface is translated into a list of coordinates that define the tool path. The interpolation uses a cubic polynomial defined by the position and derivative of the surfaces on either side of the interpolated portion of the path.

Producing a smooth surface is critical for reducing transient motion in the lathe that results in form errors in the final part. Figure 16.1 (c) shows one rotation of the cut depth in polar coordinates, along with the first derivative with respect to θ. This shows that the velocity of the lathe's x-axis is continuous throughout the cut. The result of each mirror surface having a different radius of curvature can be seen from the slight deviation from 4-fold rotational symmetry. Figure 16.1 (d) shows an interference fringe formed between a resonator coupling surface and a coupling prism that reveals a form error due to transient motion in the lathe. The resonator surface should be nearly flat with a slight convex curvature but instead has a wavy pattern in the tangential plane. This can be significantly reduced by slowing the cut velocity, as was done for resonators subsequent to this one.

16.2 Mode structure

The principle advantage of discrete monolithic resonators over whispering gallery resonators is that the mode structure can be engineered in the same way that it can be for any optical cavity consisting of spherical or toroidal mirrors. These resonators support Hermite-Gauss modes of the form presented in section 3.4.2. Because the steady-state solutions to this type of resonator are well known and have properties that are easily calculated, a resonator with particular properties can be fabricated. The modes within a system of toroidal lenses and mirrors are calculated using ABCD matrices.

We calculate the eigenmode beam parameter at the location of the coupling surface by constructing an ABCD matrix for the resonator starting immediately after reflection from that surface. For a resonator with a width, D, the ABCD matrix is

$$\prod_{i=1}^{4} \begin{pmatrix} 1 & D/\sqrt{2} \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ -2/R_e^{(i)} & 1 \end{pmatrix} \tag{16.4}$$

where $R_e^{(i)}(\theta) = R^{(i)} \cos\theta$ for the tangential plane and $R_e^{(i)}(\theta) = R^{(i)}/\cos\theta$ for the tangential plane.

Inserting the radii for each of the mirror surfaces in eq. (16.4) gives the complete ABCD matrix for either the sagittal or tangential planes. The complex beam parameter for the

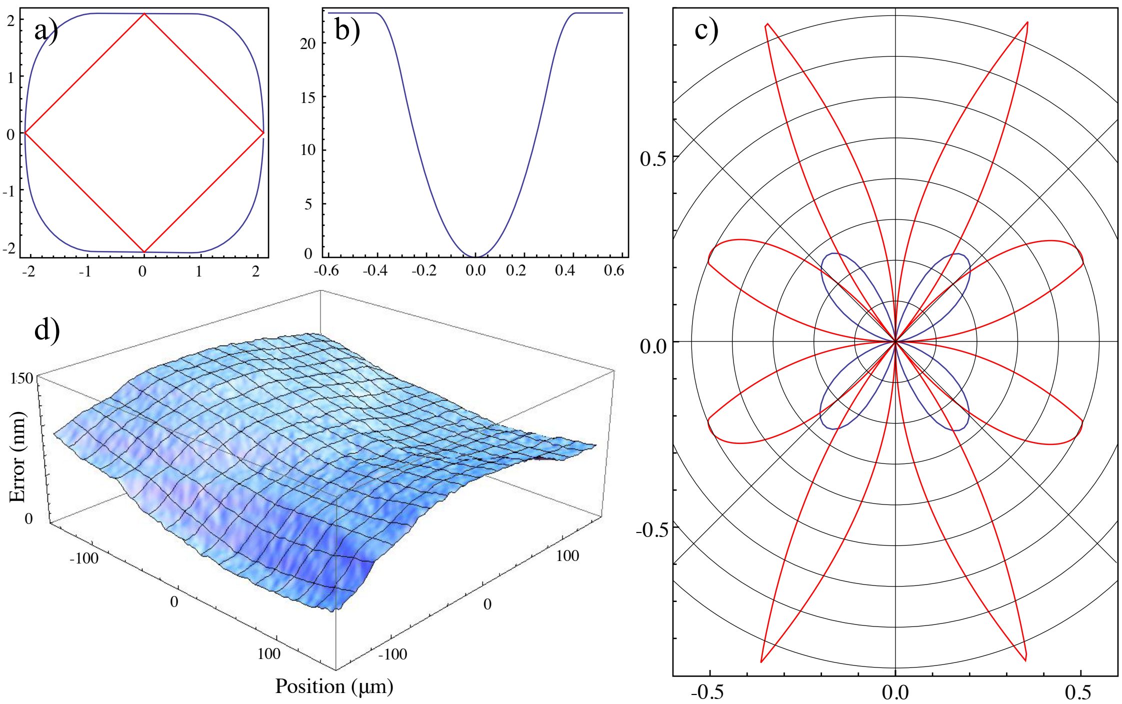

Figure 16.1: a) The tool path plus the radius of the resonator shows the targeted cut shape (blue). The red lines show the intended optical beam path within the resonator. b) The side profile of the resonator. c) The tool path for one rotation of the spindle in polar coordinates (blue) plotted with the first derivative (red) to show the tool position and velocity during the cut. d) A white-light interferometer reconstruction of a coupling mirror. The surface should be convex in both directions, but is concave in the tangential plane due to transient motion in the lathe. Later resonators were cut at a lower rotation rate to reduce this effect.

eigenmode in each plane can be be obtained by solving eq. (4.7). The mode inside the resonator can be then be engineered by tailoring the mirror radii to achieve the required mode properties, such as a non-astigmatic beam, a small mode volume or high non-linear interaction strength, as will be discussed in section 4.1.2.

The design constraints used were usually:

1. Reflectional symmetry about a line joining the two coupling surfaces and about one joining the two non-coupling surfaces. This ensures beam waists at each of the coupling surfaces and simplifies the design.
2. Nearly flat coupling surfaces. Along with focal points at the coupling surfaces, this ensures that the evanescent coupling is roughly spatially homogeneous. The curvature of the coupling surfaces was made as large as possible while still guaranteeing that the surfaces would be convex even with machining defects.
3. A round mode at some fixed point away from the coupler. This is to eliminate the need for astigmatic optics in the path of the excitation beam to obtain high mode-overlap between the excitation beam and the out-coupled light.
4. The best non-linear interaction strength within the previous constraints.

The third constraint was met by propagating the complex beam parameter of the resonator eigenmode through ABCD matrices for the coupling prism. This includes two dielectric interfaces (see table 4.1), corresponding to the resonator to prism interface and the prism to air interface, and two segments of free-space propagation, corresponding to the path inside the prism and free propagation to some fixed point. Usually, a 200 mm total path-length was used so the 200 mm focal length lenses could be used to focus the excitation beam. Because the propagation through the dielectric interfaces is not at an angle normal to the surfaces some beam distortion occurs, which is accounted for by using a slightly astigmatic circulating mode.

16.3 Prism coupling for discrete reflection resonators

The principal aim of developing the discrete total internal reflection resonators was to gain access to a predictable spatial mode structure. This is facilitated both by the Hermite-Gauss modes that are intrinsic to the resonator and by prism coupling that is spatially uniform. The latter of these qualities is a result of being able to fabricate nearly flat coupling interfaces on the resonator so that the modes do not undergo a spatial transformation as a result of the evanescent coupling. For this, we require that the shape of the resonator on the coupling interface results in a negligible change in the evanescent coupling gap across the diameter of the resonating mode as it reflects off that interface. This is a condition that is met due to the large radius of curvature of the coupling interface and because the resonating mode can be focused onto it.

Under the assumption that the coupling is uniform across the mode, the output modes will be Hermite-Gauss modes. We can mode-match the excitation laser, which is a TEM_{00} mode, to the fundamental spatial mode of the resonator using astigmatic optics. The coupling between the cavity mode and the free-space mode can then be adjusted according to eq. (14.16), using the prism distance. If the resonator has been cut to produce a circular output mode, the coupling is even simpler, requiring only standard spherical mirrors to obtain a high coupling efficiency.

Figure 16.2 shows the results of prism coupling to such a resonator using a round, TEM_{00} excitation laser. The blue traces are experimental data and the red traces are theoretical traces from eq. (14.16). The only fit parameter used in the theoretical model is the calibration of the voltage-to-distance constant for the piezo distance actuator. The data shows a 94% power coupling efficiency and is taken using a 1550 nm laser, exciting a square CaF_2 resonator that has been shaped to produce a non-astigmatic beam.

16.4 Mechanical setup

The mechanical setup of the experiment involved a number of trade-offs between design considerations and evolved throughout the experiment. The basic requirements are: excellent mechanical stability, precise adjustment of the rotation and tilt of the coupling prisms, precise and dynamic adjustment of the prism distance from the resonator, temperature control and stability of the resonator, access to apply a tuning voltage to the resonator and finally optical access for the coupling beams. The design that was settled on was to use piezo-actuated flexure mounts that were attached to compact prism mounts from Radiant Dyes. These were in turn mounted on high-stability manual translation stages from Newport. The translation stages were attached to the base. A schematic of the layout is shown in fig. 16.3 although the piezo-actuators and Peltier elements are not drawn.

The base has a pedestal in the middle, onto which the resonator is mounted. The

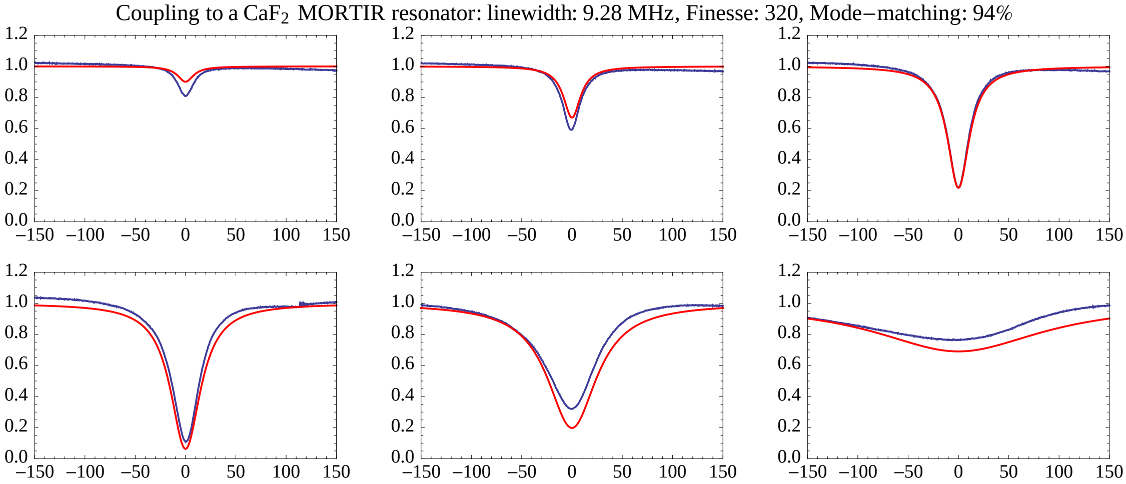

Figure 16.2: Coupling (reflection) of 1550 nm light to CaF$_2$ resonator as a function of prism (SF11) distance. The blue trace is experimental data and the red trace is from theory. The resonance is observed to pass through critical coupling and into the over-coupled regime.

thermoelectric element used to control the temperature of the resonator is located between the pedestal and resonator. Figure 16.4 (a) shows a microscope image of an early prototype resonator mounted between two prisms. The prism on the right is evanescently coupled while the prism on the left is far enough removed that no coupling occurs.

16.5 Alignment

The alignment of both the coupling prism relative to the resonator and the excitation beam are essential for observing the cavity resonance. Because the parameter space that includes both the prism orientation and the beam alignment is too large to be effectively explored, the prism must be aligned by some other means. Coarse alignment is done by directly observing the relative positions of the resonator and coupling prism using the video microscope described in appendix A.1.

Imaging the resonator and prisms from the top and then from the side allows the prisms to be rotated and tilted to match the angle of the resonator. Once the angles are correct to within the precision afforded by the microscope, the distance between the prisms and resonator is reduced to a few tens of microns, again using the microscope to monitor the relative distance. Figure 16.4 (a) shows an image taken during the process of adjusting the tilt angle of one of the prisms. The reflection of the resonator on the prism surface provides an excellent guide to fine-tune the tilt angle.

Once the coarse adjustment is complete, the alignment of the prisms can be further refined by using the Newton's rings interference fringes that form between the prisms and the resonator's coupling surface. The fringes are formed under low-coherence, monochromatic illumination. This was first provided by a sodium lamp and later upgraded to a fibre-coupled LED illuminator, described in appendix A.1.5. The illuminator had a coherence length of $\approx 10\ \mu$m, long enough that fringes could be seen while a gap was still visible between the prism and resonator on the video microscope.

Once a prism has been aligned such that the Newton's rings formed concentric circles about the centre of the coupling surface on the resonator (or uniform fringes in the case of flat coupling surfaces) the prism was aligned well enough to obtain evanescent coupling.

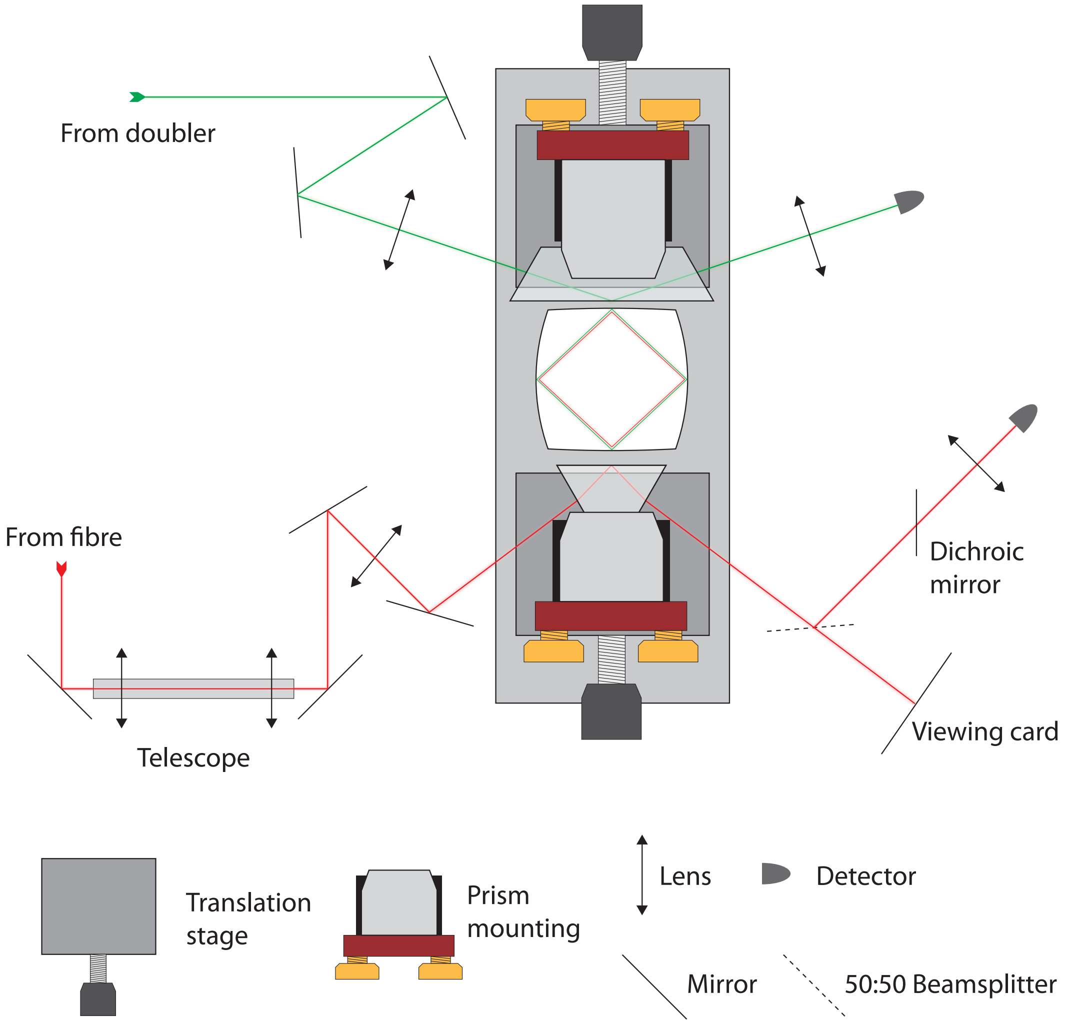

Figure 16.3: Layout of the optics for probing MOTIR resonators. The piezo-electric distance actuators are not shown.

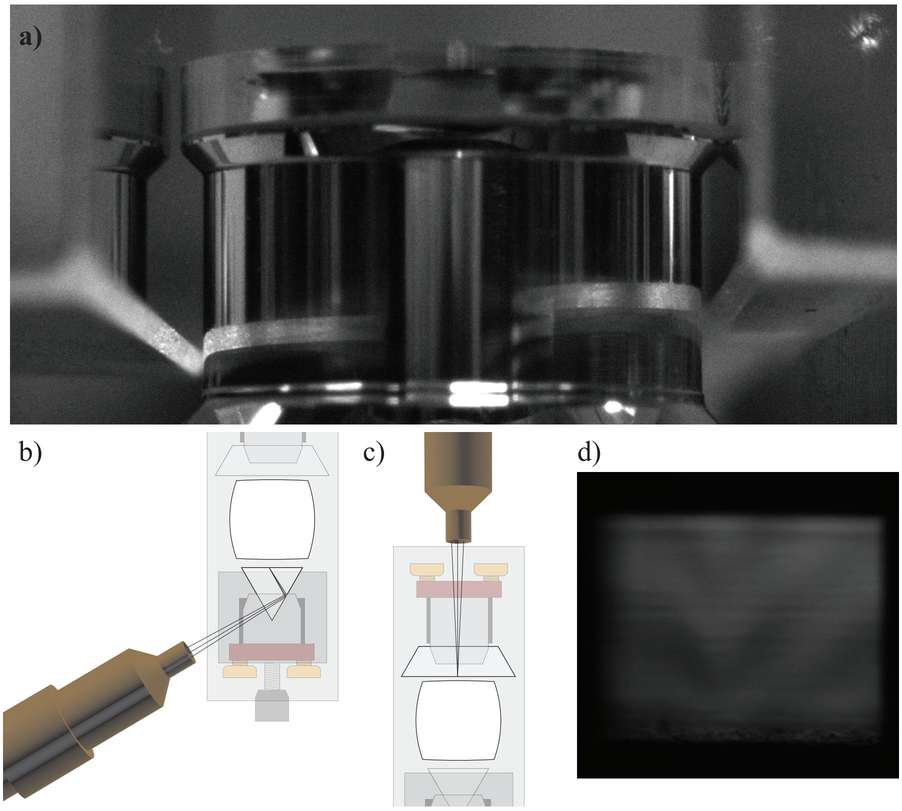

Figure 16.4: a) A resonator located between two coupling prisms. This is an early prototype and lacks the brass top electrode used for electro-optic tuning. The alignment of the video microscope used to view Newton's rings is illustrated for the SF11 prism (b) and the calcite prisms (c). An example of one of the fringe patterns observed is shown in (d). This prism has an incorrect tilt angle relative to the resonator.

Figure 16.4 (b,c) shows orientations for the microscope relative to the coupling setup for the SF11 prism and for one of the calcite coupling prisms. The rear surface of the calcite prisms was hand-polished to enable imaging of the Newton's rings through the prism while the coupling surface of the SF11 prisms could be imaged using one total internal reflection. The illumination (not shown) was usually aligned to be transmitted through the resonator and into the microscope.

This method reliably provided enough coupling between the prisms and the resonator to over-couple the cavity resonances. For very large over-coupling of the 532 nm light fine-tuning was needed after the excitation beam was aligned. This was done by carefully and iteratively changing the rotation and tilt of the prism at a safe distance from the resonator, then advancing the prism until coupling was observed and repeating the process.

16.6 Passive results

Before discussing the design of $LiNbO_3$ resonators for non-linear effects I will present the results for the discrete total internal reflection resonators acting as simple optical cavities. These demonstrate that the resonators can be machined precisely enough to have the designed mode size and shape as well as having the potential to be high-finesse passive cavities. These could find application as filter cavities, which require high-stability and low-insertion loss for TEM modes. The CaF_2 resonators are also being considered for opto-mechanical projects that couple a nano-wire oscillator to the evanescent field of the cavity. The CaF_2 resonators are ideal for this due to the spatial extent of the evanescent field because the internal angle of reflection is very close to the critical angle.

16.6.1 CaF_2

A square resonator was machined from CaF_2 , with care taken to achieve a nearly ductile mode cut. The resonator was then hand-polished to remove small defects from the machining. High-efficiency coupling into a CaF_2 resonator was obtained, as shown in fig. 16.2. The 94% efficient power coupling shows the efficacy of engineering the resonator mode to account for astigmatism in the coupling optics. The results were obtained using a round excitation beam with a waist (48 μm) roughly equal to the engineered waist (40 μm) although the coupling was not exhaustively optimised.

A relatively narrow linewidth of 9 MHz was obtained. This was somewhat surprising because the resonator was intended as a prototype and little effort had been made to polish the resonator to a high quality surface finish. This holds promise for very narrow-linewidth CaF_2 resonators in the future. That the resonator worked at all is also a demonstration of the precision of the fabrication process. The critical angle for CaF_2 is 44.52°, indicating that the resonating mode is square to within 0.5°.

16.6.2 $LiNbO_3$

The principal aim of the resonator project was to produce stable squeezing from a monolithic resonator. To achieve this, the resonator must be produced from a material that has an electro-optic effect and that can be phase-matched using birefringence to compensate for dispersion between the fundamental and second harmonic. $LiNbO_3$ satisfies these requirements for a fundamental wavelength around 1064 nm, which coincides with the wavelength available in stable monolithic Nd:YAG lasers (see section 14.1).

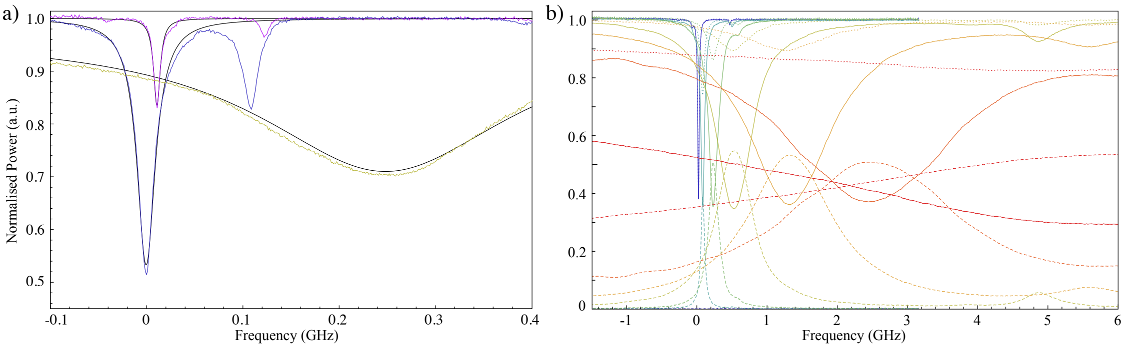

Figure 16.5: Resonances of the $LiNbO_3$ excited with 1064 nm light. In (a), only one prism is used to demonstrate under-coupling, critical coupling and over-coupling by varying the prism distance. Lorentzian fits are show in in black. (b) shows results taken with two prisms that are adjusted such that impedance matching is maintained through various coupling rates. The solid lines are the intensity reflected from the input prism and the dashed lines show the intensity transmitted through the second prism. The dotted line shows the intensity that strikes neither detector and can be assumed to be the loss in the resonator.

Before we attempted to demonstrate non-linear effects in $LiNbO_3$, the passive performance of the resonators and associated coupling optics were characterised. The most fully developed resonator, which had the best form accuracy, was polished to a high finesse, and had a top electrode attached, was a resonator that was 2.22 mm in diameter and 390 μm thick.

For 1064 nm light, the performance of the resonator was characterised using a widely tunable diode laser (see section 14.1) because of the faster scan rate than the Nd:YAG laser. Figure 16.5 shows a resonance for one of the $LiNbO_3$ resonators at three prism distances along with Lorentzian fits. The linewidths are calculated to be (8.4 ± 0.2) MHz, (23.3 ± 0.3) MHz and (322 ± 2) MHz, and correspond to under-coupling, critically coupling, and over-coupling the resonator, respectively.

The intrinsic linewidth of (8.4 ± 0.2) MHz is consistent with minimum linewidths observed for whispering gallery resonators measure in section 15.5 and the linewidth of 8.3 MHz reported in [a13]. This suggests that the observed linewidth is limited by the bulk material absorption of $LiNbO_3$ as opposed to scattering off the mirror surfaces. The strong over-coupling that can be seen in fig. 16.5 is of particular importance because it indicates that a high escape efficiency is achievable for the fundamental.

If multiple prisms are used to couple a single resonator, the distance of one of the prisms can be tuned so that the coupler is impedance matched to the loss due to the intrinsic cavity linewidth plus the loss due to the other coupler. If the loaded linewidth due to the couplers is large compared to the intrinsic linewidth, loss in the system is negligible and light can be efficiently coupled from one coupler to the other via the resonator.

Figure 16.5 demonstrates this. The resonator is evanescently coupled to two prisms, one calcite and one SF11. The distance of the two prisms is adjusted such that the input is impedance-matched to the total loss of the resonator and the second prism. The impedance-matched linewidth increases from 23.3 MHz in the absence of an out-coupling prism to more than the free-spectral range of the resonator, 21.2 GHz (not shown). This represents a completely tunable finesse ranging from 2500 (under-coupled) down to less than 1. The maximum visibility contrast shown here, 65%, is well below the best achieved for this resonator (87%) and was a result of imperfect alignment as opposed to a fundamental limit.

16.7 Independent coupling for the fundamental and second harmonic

The fundamental and second harmonic optical fields can be independently coupled to the resonator by making use of birefringent prisms made from calcite. The refractive index of calcite for 532 nm and 1064 nm is summarised in appendix B.3.2 and we can use this to engineer coupling prisms that are polarisation specific. Any coupling prism with a refractive index higher than $n_{LN} \sin 45° = 1.577$ will result in frustrated total internal reflection. Conversely, a prism with a lower index will not result in any coupling, even for very small prism distances.

Appendix B.3.2 tells us that light polarised along the ordinary axis in calcite will couple to the resonator. A 1064 nm coupler is easily engineered by orienting the optical axis parallel to that of the $\mathrm{LiNbO_3}$ resonator such that the extraordinarily polarised 532 nm light does not couple but the ordinarily polarised 1064 nm light does. Orienting the optical axis of the calcite in the plane of propagation results in a prism that will couple 532 nm light, however, the angle of the optical axis within the plane of propagation must be chosen carefully to avoid coupling the 1064 nm light as well.

Light that is polarised at some angle between ordinarily polarised and extraordinarily polarised will see a refractive index given by

$$n(\theta) = \left(\frac{\sin^2 \theta}{n_e^2} + \frac{\cos^2 \theta}{n_o^2} \right)^{-\frac{1}{2}} . \tag{16.5}$$

Snell's law for light passing from the resonator into the coupler will then be

$$n_{LN} \sin 45° = n_p(\theta - \alpha) \sin \theta, \tag{16.6}$$

where α is the angle made between the optical axis and the vector normal to the prism base. If eq. (16.6) has no real-valued solution, then the light is totally internally reflected. A simple choice of optical axis that prevents evanescent coupling would be one that is normal to the prism base; however, this would align the prism base with a cleavage plane in the calcite and make polishing the prism difficult [c46]. Instead, an angle $\alpha = 40°$ is chosen, which prevents coupling and can be polished to an optical surface.

Figure 16.6 (a-d) show the coupling prisms used along with the angles for the out-coupled light. The polarisation specific prisms (a,b) are designed such that the exiting beam forms a 5° angle relative to the vector normal to the prism side. This ensures that the back-reflection does not couple to the resonator and drive a counter-propagating mode. Another design for a 1064 nm coupler was made (c) such that the coupled 1064 nm light strikes the prism side at Brewster's angle to eliminate loss.

In addition to the polarisation specific prisms, an off-the-shelf equilateral SF11-glass prism was used to couple both polarisations. This has the advantages of being able to probe both the second-harmonic and fundamental resonances while only having to align on prism relative to the resonator. Usually, the experiment was set up with a the SF11 prism as one coupler and then a polarisation-specific coupler on the opposing side.

Figure 16.6 (e,f) show the results from characterising the polarisation-specific couplers in this configuration. The linewidth of the resonance which is strongly coupled by the

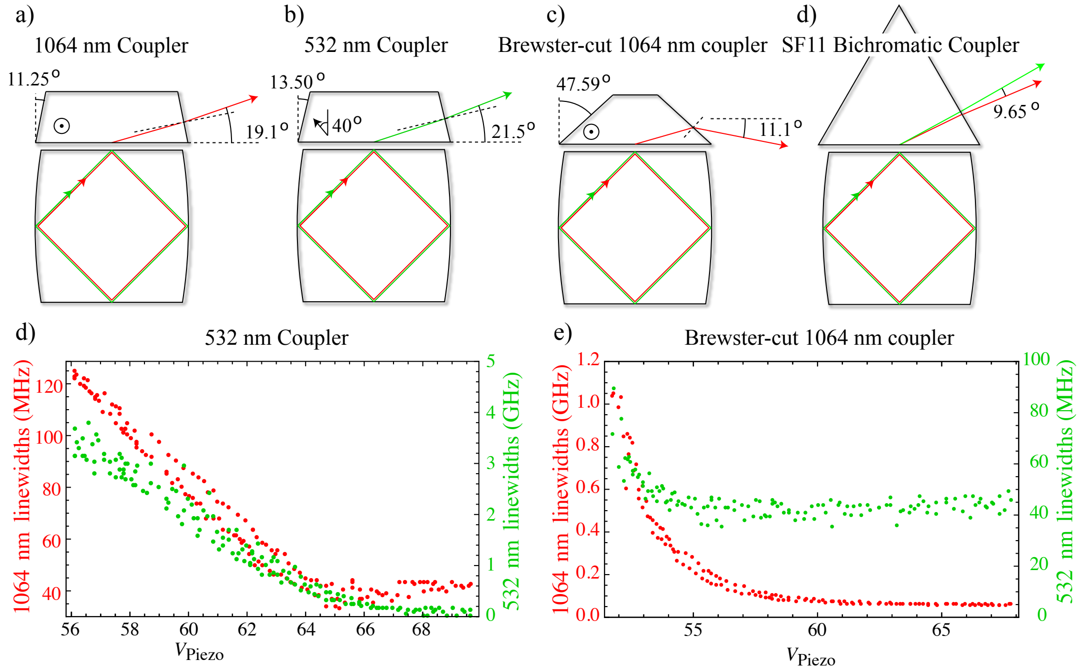

Figure 16.6: (a-d) The various prisms used for coupling the monolithic resonators. (a) and (b) are calcite prisms that couple tangentially and sagittally polarised light, respectively. (c) is similar to (a) but features Brewster's angle sides and (d) is an off-the-shelf SF11 equilateral prism. (d) Shows the linewidths for tangentially polarised 1064 nm light (red) and sagittally polarised 532 nm (light) as the piezo tuning for the 532 nm coupler is varied.

calcite prism is measured by an excitation beam that is incident on that prism. The linewidth of the resonance for the orthogonal polarisation, which should not be coupled by the calcite, is monitored using an excitation beam incident on an SF11 prism on the opposite side of the resonator from the calcite coupler.

In fig. 16.6 (e), the coupler for the vertically polarised 532 nm light is tested. The SF11 prism distance is adjusted such that the 532 nm mode is critically coupled. This over-couples the 1064 nm mode. As the calcite coupler is advanced towards the resonator by increasing the voltage to the piezo, the linewidths for both modes can be seen to increase. The ratio between the linewidths of the 532 nm and 1064 nm modes can be interpreted as the effective extinction ratio of the coupler. This was found to be 28:1. While this is not a particularly impressive figure, it should be noted that the larger extent of the evanescent field for the 1064 nm mode works against a good extinction ratio.

Figure 16.6 (f) shows the results for the horizontal coupler. Here, the linewidth of the vertically polarised 532 nm mode only begins to increase when the 1064 nm mode is very strongly over-coupled, however, the inferred extinction ratio is only 13. For both prisms it was hard to eliminate the possibility that a slight misalignment of the calcite prism caused the prism to push against the resonator and decrease the distance between the resonator and the SF11 prism, resulting in a larger linewidth than is limited by coupling to the calcite prism. For applications that require a very low finesse on one mode and a high finesse on the other, the cause of the low extinction would need to be determined.

16.8 Electro-optic tuning

$LiNbO_3$ is a linear electro-optic material: it has an index of refraction that depends on an applied constant (or radio-frequency) electric field. This is closely related to the $\chi^{(2)}$ non-linearity but the coefficients are specified for frequencies of the electric field that are significantly lower than optical frequencies. In this case the optical field responds to refractive index changes that are induced by an externally applied electric field rather than by the optical field itself. The refractive index is then $n(E) = n - \frac{1}{2}rn^3E$, where r is the electro-optic coefficient.

The optical phase shift induced by a voltage applied across an electro-optic material is then [b40]

$$\Delta_\phi = \pi V \frac{L}{d}\frac{rn^3}{\lambda_0} = \pi\frac{V}{V_\pi}, \tag{16.7}$$

where

$$V_\pi = \frac{d}{L}\frac{\lambda_0}{rn^3} \tag{16.8}$$

is the half-wave voltage.

The refractive index for the second harmonic wave responds to an electric field applied normal to the plane of the resonator according to the (unclamped) r_{33} coefficient of 33 pm/V (appendix B.1.2). The half-wave voltage for 532 nm light in the resonator presented here is then 90 V. The fundamental wave, polarised in the plane defined by the ordinary optical axes, responds according to the r_{22} electro-optic coefficient of 7 pm/V and has a half-wave voltage of 850 V.

Electro-optic tuning of the resonators is accomplished by gluing a conductive plate to the $LiNbO_3$ wafer prior to machining it into a resonator. This electrode is machined down along with the resonator so that it closely matches the resonator shape. This is too provide a uniform electric field in the $LiNbO_3$ when a voltage is applied to the electrode while providing a large surface on which to place an electrical lead without interfering with the coupling prisms. Figure 16.7 (a) shows the lead for high-voltage tuning in contact with the electrode on a resonator.

We can calculate the capacitance of the resonator and electrodes to find the cut-off frequency of the electro-optic response in a given electrical circuit. Taking the clamped (high-frequency) permittivity from appendix B.1.3, the capacitance for the parallel plate capacitor formed by the resonator and electrodes is 38 pF. In a 50 Ω circuit, this provides a 3 dB bandwidth of 84 MHz. The high-frequency feedback response could provide a very agile lock although we use a much lower loop bandwidth because the frequency instability in our system is primarily due to thermal effects.

The modulation bandwidth was tested by using a bias tee to add an RF modulation to the high-voltage tuning signal applied to the resonator. Figure 16.7 (b) shows the results for a 10 V peak-to-peak modulation at 80 MHz. The power in the each of the side-bands, roughly 8% of the central peak, corresponds to a frequency modulation index of ≈ 0.5. This is in rough agreement with the expected modulation depth, 0.78, calculated from the half-wave voltage of 850 V and assuming a 3 dB attenuation of the signal. A more thorough characterisation would require the transfer function of the bias tee and the mechanical resonant frequency of the resonator to be taken into account.

One interesting effect that was observed while electrically tuning some resonators that were operating them at high temperatures (60° C) was a reduction in the tuning range over time. We hypothesised that this may be a result of the high electric field causing domain

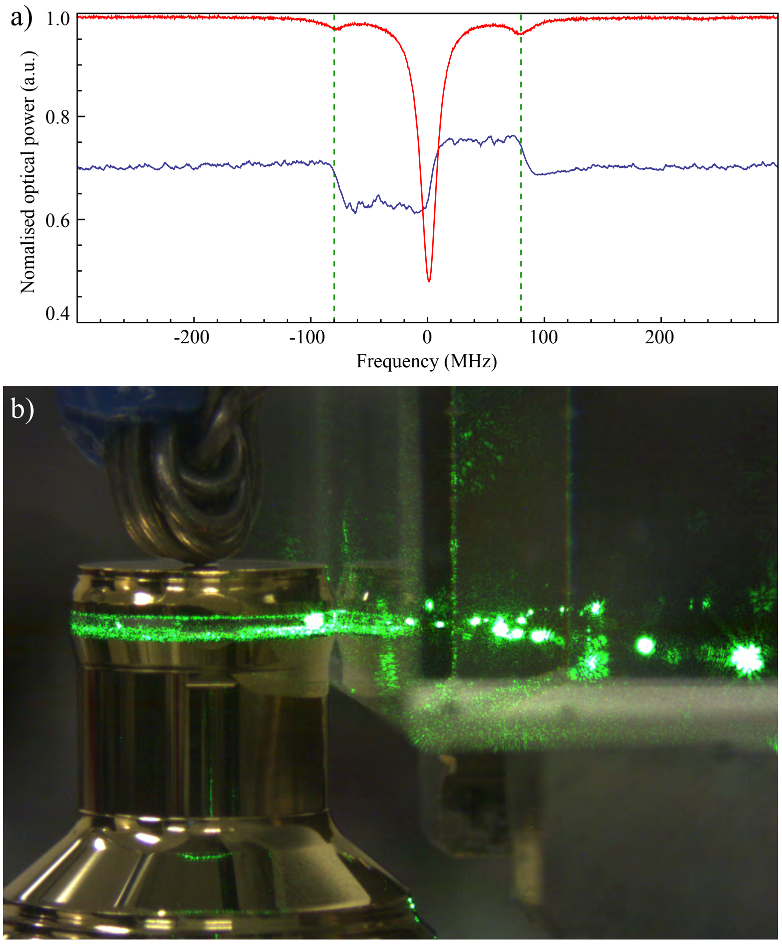

Figure 16.7: a) Modulation side-bands at 80 MHz are visible beside a cavity resonance at 1064 nm (red trace). The central frequency of the resonator itself is being modulated by the application of an RF modulation directly to the top electrode. An error signal derived from the modulation is shown in blue. b) A high-voltage lead is placed in contact with the top electrode on the resonator.

reversal in the $LiNbO_3$. To test this, we left the resonator for 15 minutes at a temperature of $\sim 120°$ C while applying 1 kV across the resonator in an attempt to completely reverse the domains. After this time, the tuning range returned to what it had been originally and didn't exhibit any further reductions. This effect, observed in two separate resonators seems to support the hypothesis but warrants further investigation particularly because, while the electric field is of a sufficient magnitude for domain reversal, the temperature is well below the Curie temperature for $LiNbO_3$ [a47].

16.9 Phase-matching

In order to predict the temperature that is needed to phase-match at a particular fundamental wavelength, the refractive indices of the ordinary and extraordinary waves need to be known as functions of wavelength and temperature. This is usually expressed by the Sellmeier equation.

In its general form, the Sellmeier equation is

$$n^2(\lambda, T) = 1 + \sum_i \frac{A_i \lambda^2}{\lambda^2 - \lambda_i(T)^2}, \tag{16.9}$$

where the sum terms represent absorption resonances, i, with strengths proportional to A_i at wavelengths λ_i. The temperature dependence is incorporated into the wavelengths of the absorption resonances. The phase matching temperature can be obtained numerically by equating the Sellmeier equation for the fundamental field with that of the second harmonic.

$LiNbO_3$ is a popular non-linear material because it can be critically phase-matched at low temperatures for 1064 nm, the wavelength emitted by Nd:YAG and Nd:YVO_4 lasers. Our lab uses monolithic non-planar ring oscillator (NPRO) Nd:YAG lasers supplied by Innolight (now part of Coherent Scientific) along with $LiNbO_3$ or PPKTP non-linear crystals to generate quantum light. The monolithic geometries are well suited to non-critical phase matching because of the changing propagation direction with the crystal, therefore, $LiNbO_3$ is the natural choice for this application. It can be non-critically phase-matched at moderate temperatures by orienting the fundamental in the plane defined by the ordinary axes and the second harmonic along the extraordinary (optical) axis.

Obtaining an accurate phase-matching temperature requires precise values for the coefficients for Sellmeier equation for the MgO:$LiNbO_3$ used in our experiment. The stoichiometric composition of the $LiNbO_3$ and the doping concentration of magnesium oxide (MgO) affect these coefficients, and thus the phase-matching temperature. Good quality control of the crystal composition is therefore crucial to predicting the optical properties. We use MgO:$LiNbO_3$ from Crystal Technology (now part of Gooch and Housego) which are known for providing excellent quality materials.

Using the stoichiometric ratios provided by the $LiNbO_3$ supplier, we use the Sellmeier equations from [a48] which are provided in terms of the composition of the $LiNbO_3$ as well as the temperature (see appendix B.1.2). The specified molar ratio of Li_2O in the composition is 43.38 ± 0.01 mol% with an MgO doping concentration of 5%. According to the equations in [a48], this would yield a phase-matching temperature of 112.25°C for 1064 nm. The phase-matching temperature specified by the supplier is 109.0°C for this wavelength, which would correspond to an MgO doping concentration of 4.85%. Using this doping concentration (4.85%), we obtain a phase-matched refractive index of 2.2286 at 1064 nm.

The Sellmeier equations also allow us to predict the phase-matching temperature for any wavelength. Using the same parameters that predict the correct phase-matching temperature for 1064 nm, we can numerically solve the Sellmeier equations for phase-matching over the wavelength range of our tunable diode laser (1030 nm - 1070 nm). Interpolating the results with a third-order polynomial yields:

$$T_{pm}(\lambda) = -6.6331 \times 10^4 + 1.7649 \times 10^2 \lambda - 1.5772 \times 10^{-1} \lambda^2 + 4.7495 \times 10^{-5} \lambda^3, \tag{16.10}$$

where λ is in nm and T_{pm} is in °C. This allows us to operate at a much lower temperature (eg. 40°C for 1035 nm) than at 1064 nm.

The phase-matching curve was measured using a square monolithic resonator with a width of 2.2 mm. Points at which the phase-matching condition was satisfied were found for a set of temperatures between 40°C and 100°C by adjusting the wavelength of the diode laser until a maximum intensity of second-harmonic was detected. The phase-matching

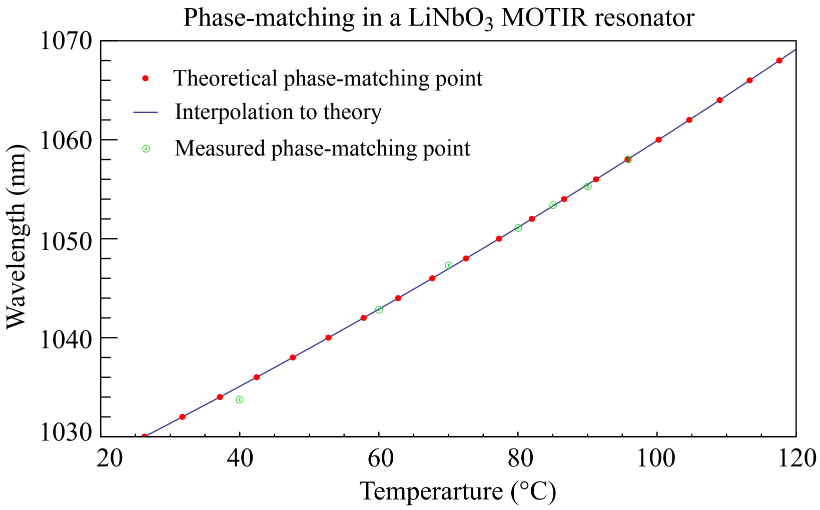

Figure 16.8: Phase-matching temperatures (red) for the MgO:LiNbO$_3$ supplied by Crystal Technology, calculated from eq. (16.10). Experimental phase matching points (green) are shown as open circles with error bars inside the circles.

condition was scanned in a small window using the voltage across the resonator. The results are shown in fig. 16.8 plotted with the predicted phase-matching curve. Error bars correspond to the tuning range of the voltage scan.

16.10 Non-linear conversion efficiency

The efficiency of a non-linear process depends on the mode volume of the interaction, the amplitude of the fields and the relative phase between them. The optimal beam size for a given crystal can be calculated as a compromise between a small mode waist and a longer Rayleigh range, complicated by the fact that the fundamental and second harmonic will accumulate a different Gouy phase shift across the beam waist. The optimal Rayleigh range is known as the Boyd-Klienman parameter. For the monolithic resonators the calculation needs to be done to include each segment of the resonator, summed with the appropriate phase shifts to account for the total internal reflections and with the requirement that both the fundamental and second harmonic are resonant in the cavity.

16.10.1 The Boyd-Kleinman condition

We can calculate the efficiency of the SHG process by making a series of simplifying assumptions which are well-justified in the cases that are of interest for this thesis. Firstly, we assume a Gaussian fundamental wave from eq. (3.38):

$$E = E_0 u_x(x,z) u_y(y,z) \tag{16.11}$$

$$u_x(x,z) \propto \tilde{q}_x^{-\frac{1}{2}}(z) e^{ik\frac{x^2}{2q_x(z)}}, \tag{16.12}$$

where $\tilde{q}(z) \equiv q(z)/z_R$ is dimensionless and the e^{ikz} term is included in E_0. Inserting this into eq. (4.3) and integrating over the volume gives

$$E_{SHG} = \frac{ik_{2\omega} d_{eff}}{2n_{2\omega}} \int_V E_0^2 \frac{z_{Rx} z_{Ry}}{q_x(z) q_y(z)} \exp 2ik \frac{x^2}{2q_x(z)} + \frac{y^2}{2q_y(z)} e^{i\Delta kz} dV. \tag{16.13}$$

We can identify the same envelope function as in eq. (16.13) for the second harmonic and simplify to

$$E_{SHG} = \frac{ik_{2\omega} d_{eff}}{2n_{2\omega}} E_0^2 \sqrt{z_{Rx} z_{Ry}} \int_V \frac{u_{(2\omega,x)} u_{(2\omega,y)}}{\sqrt{q_x(z) q_y(z)}} e^{i\Delta kz} dV. \tag{16.14}$$

The optical power in each segment is related to the electric field according to $P = 1/2c\epsilon_0 n \int |E|^2 dA$. We can perform this integral independently of z and assume that u_x and u_y are normalised such that

$$P = \frac{1}{4} c\epsilon_0 n\pi w_x w_y |E_0|^2 \tag{16.15}$$

where w_x and w_y are the beam waists for the fundamental wave in the two transverse direction. Using the same relation for optical power for the second harmonic and performing some algebra, we arrive at

$$P_{SHG} = \frac{2}{\pi} \frac{\omega^3 d_{eff}^2}{n^2 c^4 \epsilon_0} P_0^2 \left| \int_z \frac{e^{i\Delta kz}}{\sqrt{\tilde{q}_x(z) \tilde{q}_y(z)}} dz \right|^2. \tag{16.16}$$

Here, ω is the frequency of the fundamental wave and n is the refractive index, which must be common to both fields for phase-matching to occur. We can define a single-pass conversion efficiency for the process

$$\eta_{SHG} = \gamma_{SHG} L G P_0, \tag{16.17}$$

where

$$\gamma_{SHG} = \frac{2}{\pi} \frac{\omega^3 d_{eff}^2}{n^2 c^4 \epsilon_0}, \tag{16.18}$$

and L is the length of the interaction. For LiNbO$_3$ the coefficient $\gamma_{SHG} = 0.22/\mathrm{Wm}$ [a7]. The parameter

$$G = \left| \sqrt{\frac{z_R}{L}} \int h(\tau) d\tau \right|^2 \tag{16.19}$$

where

$$h(\tau) = \frac{e^{i\Delta k\tau}}{\sqrt{(1 + i\tau/\alpha)(1 + i\alpha\tau)}} \tag{16.20}$$

is the well-known Boyd-Kleinman focusing factor [a49, c50] for the case in which there is no beam walk-off. For simplicity, a mean Rayleigh length z_R has been used and a parameter α is introduced to keep track of astigmatism. For single-pass second-harmonic generation, G can be maximised by ensuring that $L = 2.84 z_R$.

16.10.2 Optimising the non-linearity

The geometry of the monolithic resonators can be optimised for non-linear performance by selecting the radius of curvature for the mirrors that provides the highest non-linear interaction between the fundamental and second harmonic waves. For single-pass second-harmonic generation, the non-linear interaction optimised by numerically maximising eq. (16.19) using the Rayleigh length, z_R and the phase matching condition Δk as free-parameters. In this simple case, Δk can be selected to compensate for the different Gouy phase shift, eqs. (3.42), between the fundamental and second harmonic.

The same principle applies for the monolithic resonators, however the situation is complicated by multiple beam waists, phase-shifts between the fundamental and second-harmonic waves at each total internal reflection and the requirement that both the fundamental and second waves must be resonant in the cavity. We follow the calculation performed in [a7] for exactly this system. Double resonance is satisfied when the following conditions are met:

$$\frac{n_o(\omega)\omega}{c}L - \phi_x - \phi_y + \sum_{j=1}^{N} \gamma_p^{(j)} = 2\pi M_0 \tag{16.21}$$

$$\frac{n_{2\omega}2\omega}{c}L - \phi_x - \phi_y + \sum_{j=1}^{N} \gamma_s^{(j)} = 2\pi M_{SHG}. \tag{16.22}$$

$$\tag{16.23}$$

L is the total path length, $n_{o,e}$ is the refractive index for the ordinary or extraordinary wave, $\phi_{x,y}$ is the total Gouy phase shift for each transverse direction for one round trip, and $\gamma_{p,s}^{(j)}$ is the phase-shift due to total internal reflection on the $j-$th mirror surface. These phase shifts can be calculated from the Fresnel reflection coefficients eqs. (14.11).

To calculate the focusing factor for the entire cavity we must stitch together the focusing factors for each leg, while taking care to add the correct phase to each. The resulting modification to 16.19 gives

$$G_{Res} = \left| \sqrt{\frac{z_R}{L}} \sum_{i=1}^{N} \left(\int h(\tau) d\tau \prod_{j=1}^{i-1} e^{i(2\gamma_p^{(j)} - \gamma_s^{(j)})} \right) \right|^2 . \tag{16.24}$$

This focusing factor can be calculated for a proposed cavity design by evaluating eq. (16.24). In practice, this was done by finding the stable solution for the complex beam parameter. The integrals in eq. (16.24) were then evaluated numerically for each segment of the cavity. The results were then summed with appropriate phase offsets to account for the total internal reflections. Because the complex beam parameter provides the Gouy phase relative to the beam waist and the cavity has four waists, care was taken to first add appropriate phase offsets to provide a continuous phase, after which the additional phases for total internal reflection were added.

Figure 16.9 shows some results from a numerical analysis of the non-linear conversion efficiency. The calculation is shown is for the resonator described in section 16.6.2 and is done by using ABCD matrices to find the eigen-mode of the resonator. Figure 16.9 (a) plots the beam waist and Gouy phase shift for one round-trip through the resonator. The magnitude and phase of eq. (16.20), the relative contribution to the non-linearity at point along the optical propagation, are shown in Figure 16.9 (b). The total magnitude

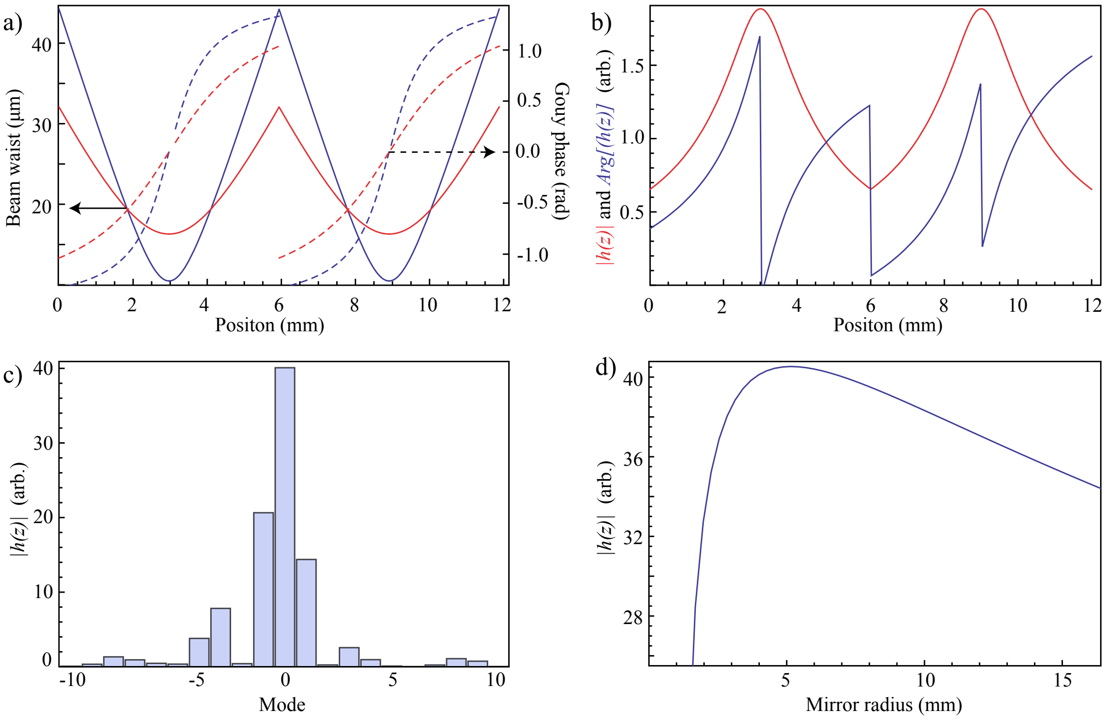

Figure 16.9: Properties of the eigen-mode of a square $\mathrm{LiNbO_3}$ resonator through one round-trip. (a) The beam waist and Gouy phase. (b) The local contributions to non-linear interaction, given by eq. (16.20). (c) An example of the non-linear conversion for subsequent longitudinal modes of the resonator, centred about the one corresponding to $\Delta_k = 0$. (d) The effective non-linear interaction strength as a function of mirror radius for a particular size of resonator.

to the effective non-linear coefficient is also shown as a function of the longitudinal mode number relative to the mode number with the highest non-linear coefficient in (c) and as a function of the mirror curvature for the mode with the maximum coefficient (d). Adjusting the mirror radius until the coefficient is at a maximum allows optimisation of the strength of non-linear interaction in the resonator.

16.11 Photo-refractive effects

The square monolithic resonators proved to be successful in overcoming the challenges associated with coupling the whispering gallery resonators. Excellent control over the coupling rates was demonstrated as well as proof-of-concept non-linear conversion. Ultimately, however, increasing the circulating power of the second harmonic in the resonator, whether as a result of second harmonic generation in the resonator or as a pump for down-conversion, led to instability due to photo-refractive effects. These are related to absorption of visible light, and are only significant with the 532 nm second harmonic field.

The photo-refractive effects arise from iron impurities in the crystal [a51–a53]. These impurities result in a photo-voltaic effect that, in turn, affects the refractive index through the electro-optic effect. The first approach taken to reduce photo-refractive effects was therefore to attempt to source crystal with the highest purity.

$\mathrm{LiNbO_3}$ was obtained from United Crystals, Gooch & Housego, and from a sample that had been previously purchased from Casix for use in squeezing experiments. All of the $LiNbO_3$ was magnesium-doped $LiNbO_3$ (MgO:$LiNbO_3$) because the optical damage threshold is increased and photo-refractive effects are reduced [a54]. Additionally, the phase-matching temperature in MgO-doped $LiNbO_3$ is significantly lower than for the undoped material (see appendix B.1.2). The $LiNbO_3$ samples from United Crystals had doping concentrations of 3%, 5% and 7% while the samples from Gooch & Housego and Casix were 6.5% MgO-doped. Of these, the Casix sample exhibited the least photo-refractive effects while the United Crystals exhibited the most.

The mechanism behind the reduced photo-refractive effects in MgO-doped $LiNbO_3$ is an increase in photo-conductivity [a54]. This additional conductivity under visible illumination allows optically induced charges in the crystal to be neutralised, reducing the photo-voltaic effect. The mechanism, however, is problematic when combined with the large tuning voltages that we applied to frequency-lock and phase-match the resonators. The photo-conductivity was observed to have the effect of a high-pass filter on the frequency tuning. The resonator would respond to rapid changes in tuning voltage but photo-conductivity would quickly compensate for the applied potential. The magnitude of the effect is related to the optical power of 532 nm light. The combination of voltage tuning with 532 nm light was therefore observed to cause significant photo-refractive effects in the crystal.

The photo-refractive effects are considerably more problematic when the second harmonic mode has a high-finesse. In this case, the intra-cavity field is more intense and small fluctuations in the refractive index will have a more significant effect relative to the linewidth of the resonance. We therefore attempted to use independent coupling of the second harmonic and fundamental, section 16.7, to reduce the finesse of only the second harmonic. Additionally, we used a split-feedback to the resonator, using temperature of the resonator for slow feedback to the frequency and the tuning voltage for fast feedback. This kept the tuning voltage low to reduce the magnitude of photo-refractive effects.

These measures were a partial success. We were able to briefly lock the resonator frequency to a bright second-harmonic pump and observe parametric amplification of the fundamental. Unfortunately, the lock was very unstable. Careful improvements to the locking system would be likely to render the system stable enough to observe squeezed light. The reduction in the impact that the photo-refractive effect has on the system, however, comes at the expense of voltage tunability and low pump power operation due to a low-finesse for the second harmonic. With these trade-offs, the added complexity of the evanescent coupling seems to have few advantages over a more traditional cavity design with dielectric coatings.

CHAPTER 17

Resonator-based memory

The compact nature of monolithic resonators combined with the long photon lifetimes that have been demonstrated with them offers the possibility of using them for simple quantum memories for light. A pair of coupled resonators that are in turn coupled to a propagating mode, as shown in figure 17.1, provide a platform for exploiting the favourable properties of tunable monolithic resonators for use as a memory. In this design, a resonator with a long photon lifetime is coupled to a propagating optical mode only through a tunable intermediate resonator. This intermediate resonator acts as a tunable output coupler for the resonator with the longer lifetime.

The combined system of the two optical resonators acts as a Q-switched resonator. Provided that light can be efficiently coupled into the system, the quality factor can be dynamically increased and the light will be held within the storage resonator until the quality factor is again reduced and the light is released. The system can be modelled using coupled mode theory. We consider a tunable resonator with a circulating mode a with an intrinsic decay rate γ_1 that is coupled to a travelling mode with with a coupling rate Γ. This resonator is further coupled to a second resonator with circulating mode b with a coupling rate κ. This second resonator decays with an intrinsic decay rate γ_2. The coupled equations of motion for the system are

$$\partial_t a(t) = \left(i\omega(t) - \Gamma - \gamma_1\right) a(t) + i\kappa b(t) + \sqrt{2\Gamma} s_{in}(t) \tag{17.1a}$$

$$\partial_t b(t) = -\gamma_2 b(t) + i\kappa a(t) \tag{17.1b}$$

$$s_{out}(t) = s_{in}(t) + \sqrt{2\Gamma} a(t), \tag{17.1c}$$

where we have assumed that the mode b is resonant with the incoming light and that a is detuned by $\omega(t)$. The travelling mode is driven by an incoming field s_{in} and s_{out} is the output.

The equations of motion bear a resemblance to those governing a three-level atom. Here, the circulating modes play the role of the atomic coherences and the coupling rates play the role of the light-atom coupling rate. Similarly to our treatment of the driven three-level atom, the intermediate mode, in this case a, can be adiabatically eliminated by assuming that the coupling rate Γ is large relative to the rate at which a varies. The implication is that the loaded linewidth of the intermediate resonator must be large relative to the bandwidth of the pulses that are being stored. Making the adiabatic approximation

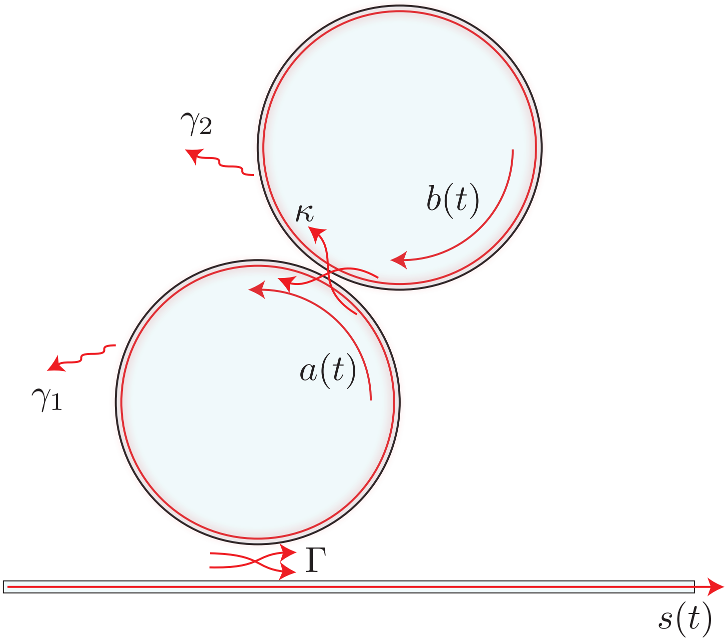

Figure 17.1:

that $a'(t) \approx 0$ the equations of motion can be written as

$$\partial_t b = -\left(\gamma_2 + \beta\Omega\right) b - \sqrt{2\beta}\Omega s_{in} \tag{17.2}$$

$$s_{out} = \left(1 - 2\Omega\right) s_{in} + \sqrt{2\beta}\Omega b \tag{17.3}$$

with the definitions

$$\beta \equiv \kappa^2/\Gamma; \quad \Omega \equiv \cos\theta e^{i\theta}; \quad \tan\left(\theta\right) \equiv \omega/\Gamma. \tag{17.4}$$

The function $\Omega(t)$ depends on the detuning of the intermediate resonator and acts as a Q-switch for b. In the limit of $\omega \to \infty$, Ω asymptotically approaches zero and the quality factor of b becomes dominated by the intrinsic decay rate γ_2. When a is resonant with b, however, $\Omega = 1$ and b is directly coupled to the travelling mode with coupling rate β.

With arbitrary control over the coupling to b it is possible to trap a circulating mode until it is needed and then release it. In order to use the resonators as a memory, however, it is necessary to efficiently couple light into the mode b. This can be accomplished by tailoring the input pulse to be a time-reversed version of the pulse that would result at the output by allowing the resonator mode to decay. If the quality factor of the storage resonator can be switched instantaneously, the output pulse would have an exponential decay profile with a time constant β. The time-reversal of this would therefore be an exponentially rising ramp.

Such a pulse shape, however, is impractical for a memory. Incoming photons would have to be shaped into the ideal temporal mode and the output from one memory would have a poor overlap with the ideal input for the next memory. The detuning of the coupling resonator however, can be used to tune the coupling rate of the strange resonator to any value between 0 and β as a function of time. It can therefore be used to temporally shape the intensity of the output pulse emerging from the storage resonator, although there will be an additional time-varying phase introduced.

17.1 Possible implementations

The general concept of two coupled resonators leaves a wide variety of possible physical implementations available. Here, I will consider variations of monolithic resonators that were produced and tested during my studies. The options for geometries are discrete total-internal reflection resonators and whispering gallery resonators. The discrete reflection resonators have the advantage of high evanescent coupling rates while the whispering gallery resonators have been demonstrated to have very long photon lifetimes. Options for materials include CaF_2 , MgF_2, $LiNbO_3$, and KTP. The requirement that the coupling resonator be rapidly tunable restricts the options to one of the electro-optic materials while the requirement of a long storage lifetime would be best satisfied with CaF_2 or MgF_2.

Two possibilities for implementations with different advantages and disadvantages are worth considering. The simpler of the two is to use the same electro-optic material for both resonators. This allows identical geometries to be used for each resonator because there is no deviation in the angle of the beam transmitted from one resonator to the next due to refraction. In this case, two square resonators or two whispering gallery resonators could be used, however the high coupling rate that is required between the coupling resonator and the prism suggests using a square discrete total internal reflection. Because KTP has less bulk absorption than $LiNbO_3$ in the frequency range of interest, it is the better candidate as a material. Figure 17.2 (a) shows the layout for this concept.

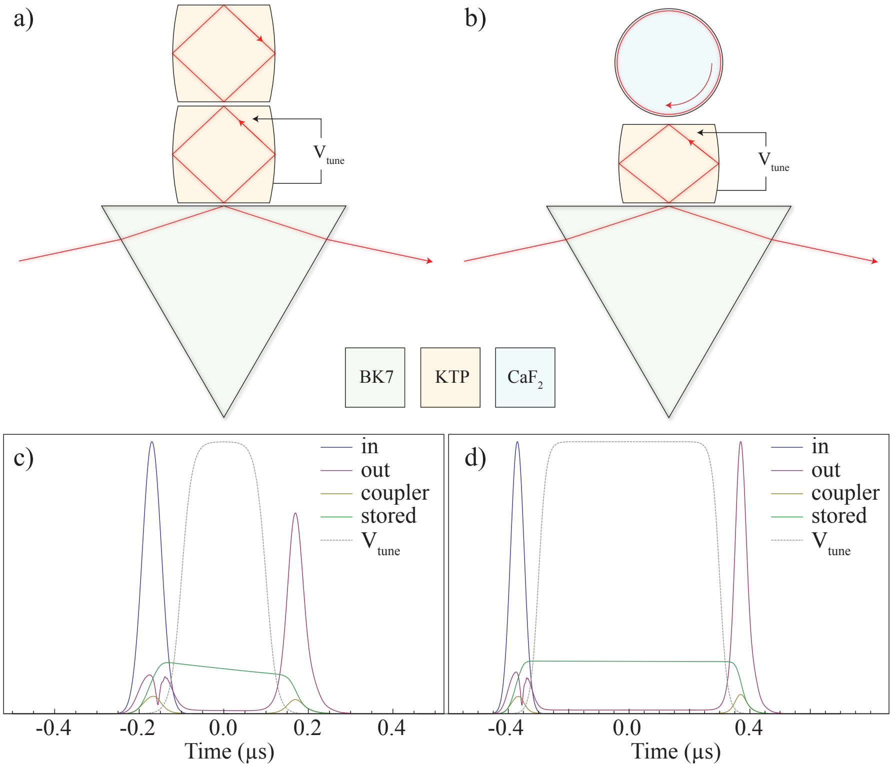

Figure 17.2: a) $\kappa = 100$, $\Gamma = 0.1$, $\gamma_1 = \gamma_2 = 1$, $0 < \omega < 2 \times 10^4$, b) $\kappa = 100$, $\Gamma = 0.1$, $\gamma_1 = \gamma_2 = 0.01$, $0 < \omega < 2 \times 10^4$

The second implementation is to use two different materials for the storage and coupling resonators to allow CaF_2 to be used as the storage resonator. In this case, the aspect ratio of the coupling resonator will have to be adjusted to phase-match with the evanescent field of the storage resonator. This allows flexibility for the geometry of the storage resonator. Whispering gallery resonators are a good candidate because of the small observed linewidths and a demonstration that they can be over-coupled, as shown in fig. 15.4.

Both of these cases were investigated by numerically solving eqs. (17.1a) to (17.1c) using optimistic but plausible parameters. The simulations assume an intrinsic loss rate of 1 MHz in KTP and 1 kHz [a3] in CaF_2 . The extremely long lifetime for CaF_2 is three orders of magnitude higher than rates measured in our experiments and was achieved using a sequence of polishing and annealing steps but has been observed in a whispering gallery resonator. The temporal shape of the tuning curve was restricted to the form of a logistic curve rather than tailored to the optimal switching profile. The reason is that high-voltage arbitrary waveforms are challenging to produce experimentally, however, a rough approximation of a logistic curve can be produced by a sharp step-function and a passive Bessel filter, both of which are straightforward to obtain.

17.2 Limitations

Aside from the limitations on memory lifetime, the memory is subject to a set of constraints on performance. The coupling resonator must have a loaded linewidth that is large relative to the time-scale of the pulses being stored while maintaining a large finesse. It must also be rapidly tunable on the same time-scale. The simulations were done for pulses with a characteristic bandwidth of 10 MHz, and assumed a loaded linewidth of 100 MHz with a loaded finesse of 200 for the coupling resonator. A KTP resonator similar to the $LiNbO_3$ resonator discussed in chapter 16 would satisfy these requirements.

Figure 17.2 (c) shows the results of numerically solving eqs. (17.1) using the literature values of KTP in geometry similar to those of the $LiNbO_3$ resonators that we produced and worked with. The required coupling rates would be easily achievable, as would the required tuning bandwidth. While alignment of the two resonators would be challenging, this challenge could likely be met using the interferometric techniques discussed in section 16.5. The memory lifetime, however, would be limited by the bulk absorption in KTP, as shown in Figure 17.2.

A CaF_2 whispering gallery resonator would comfortably satisfy the requirements for coupling rate. As evidenced by fig. 15.4 the coupling rate could provide storage for photons with at least a 9 MHz bandwidth and could have a storage life time of up to 1 ms [a3]. Alignment would also be simplified by the cylindrical symmetry of the whispering gallery resonator. The challenge would rest in producing a rectangular monolithic $LiNbO_3$ resonator with an internal angle of reflection that is appropriate for coupling to the CaF_2 resonator.

CHAPTER 18

Conclusions and outlook

Ultra-precision machining was demonstrated to be an effective method for producing high quality factor monolithic resonators from crystalline optical materials. While the technique had previously been applied to the fabrication of whispering gallery resonators [a3], we have shown that it can also be extended to other shapes. By ensuring that the machining is done in the ductile regime of material removal, only minimal hand polishing is required to produce a high-quality optical finish.

The flexible machining process allows resonators to be produced with optical surfaces that feature different radii of curvature for each surface and in both the sagittal and tangential planes. This allows us to engineer the resonant optical mode to couple efficiently to circular free-space Hermite-Gaussian modes and to optimise the effective non-linear coupling coefficient.

The square monolithic resonators have also been shown to be a viable path to tunable finesse optical cavities. The coupling rate can be tuned to give a linewidth ranging from 9 MHz to more than the 21.2 GHz free spectral range of the resonator. For s-polarised light, the electro-optic effect was shown to provide a frequency tuning of more than a free-spectral range with only a 200 V potential and with a bandwidth of $\approx$ 80 MHz. The versatile tunablity may have applications in widely tunable optical filters that feature low insertion loss.

Non-linear optical frequency conversion was observed in both whispering gallery and square resonator geometries. We were able to demonstrate second-harmonic generation at low pump powers in both resonator configurations as well as observe evidence of parametric down-conversion.

In order to fully develop the resonators into compact sources of quantum light, the problem of photo-refractive effects of visible light in $LiNbO_3$ will need to be solved. We have begun work on identifying crystal samples that exhibit minimal photo-refractive effects as well as on using thermal non-linearities to provide a robust locking mechanism that can keep the resonator locked to a second-harmonic pump in spite of photo-refractive effects. Another approach would be to produce similar resonators in periodically-poled KTP (PPKTP). Such a solution would be challenging because the poling domains would need to be carefully aligned with the machined shape. The success of PPKTP as a non-linear material for the down-conversion of 532 nm light suggests that it is a path worth exploring.

Bibliography - Monolithic Resonators

Articles

a2. Matsko, A. B. & Ilchenko, V. S. Optical resonators with whispering gallery modes I: basics. *IEEE J. Sel. Top. Quantum Electron* **12** (2006).

a3. Grudinin, I., Matsko, A. & Savchenkov, A. Ultra high Q crystalline microcavities. *Optics.* <`http://linkinghub.elsevier.com/retrieve/pii/S0030401806002136`> (Jan. 2006).

a4. Risk, W. P. & Kozlovsky, W. J. Efficient generation of blue light by doubly resonant sum-frequency mixing in a monolithic KTP resonator. *Opt. Lett.* **17,** 707–709 (May 1992).

a5. Schiller, S., Yu, I. I., Fejer, M. M. & Byer, R. L. Fused-silica monolithic total-internal-reflection resonator. *Opt. Lett.* **17,** 378–380 (Mar. 1992).

a6. Fiedler, K., Schiller, S., Paschotta, R., Kürz, P. & Mlynek, J. Highly efficient frequency doubling with a doubly resonant monolithic total-internal-reflection ring resonator. *Opt. Lett.* **18,** 1786–1788 (Nov. 1993).

a7. Schiller, S. & Byer, R. L. Quadruply resonant optical parametric oscillation in a monolithic total-internal-reflection resonator. *J. Opt. Soc. Am. B* **10,** 1696–1707 (Sept. 1993).

a8. Kane, T. J. & Byer, R. L. Monolithic, unidirectional single-mode Nd:YAG ring laser. *Opt. Lett.* **10,** 65–67 (Feb. 1985).

a9. Armani, D., Kippenberg, T., Spillane, S. & Vahala, K. Ultra-high-Q toroid microcavity on a chip. *Nature* **421,** 925–928 (2003).

a10. Pöllinger, M., O'Shea, D., Warken, F. & Rauschenbeutel, A. Ultrahigh-Q tunable whispering-gallery-mode microresonator. *Physical review letters* **103,** 053901 (2009).

a11. Cai, M., Painter, O. & Vahala, K. J. Observation of critical coupling in a fiber taper to a silica-microsphere whispering-gallery mode system. *Physical review letters* **85,** 74 (2000).

a12. Gorodetsky, M. L., Savchenkov, A. A. & Ilchenko, V. S. Ultimate Q of optical microsphere resonators. *Optics Letters* **21,** 453–455 (1996).

a13. Fürst, J. *et al.* Naturally phase-matched second-harmonic generation in a whispering-gallery-mode resonator. *Physical review letters* **104,** 153901 (2010).

a14. Aoki, T. *et al.* Observation of strong coupling between one atom and a monolithic microresonator. *Nature* **443,** 671–674 (2006).

a15. Swaim, J. D., Knittel, J. & Bowen, W. P. Detection of nanoparticles with a frequency locked whispering gallery mode microresonator. *Applied Physics Letters* **102,** 183106 (2013).

a16. Schliesser, A. & Kippenberg, T. J. Cavity optomechanics with whispering-gallery mode optical micro-resonators. *Advances In Atomic, Molecular, and Optical Physics* **58,** 207–323 (2010).

a17. Bahl, G., Zehnpfennig, J., Tomes, M. & Carmon, T. Stimulated optomechanical excitation of surface acoustic waves in a microdevice. *Nature communications* **2,** 403 (2011).

a18. Hofer, J., Schließer, A. & Kippenberg, T. J. Cavity optomechanics with ultrahigh-Q crystalline microresonators. *Physical Review A* **82,** 031804 (2010).

a19. Fiore, V., Dong, C., Kuzyk, M. C. & Wang, H. Optomechanical light storage in a silica microresonator. *Physical Review A* **87,** 023812 (2013).

a21. Förtsch, M. *et al.* A versatile source of single photons for quantum information processing. *Nature communications* **4,** 1818 (2013).

a22. Kippenberg, T. J., Holzwarth, R. & Diddams, S. Microresonator-based optical frequency combs. *Science* **332,** 555–559 (2011).

a23. Del'Haye, P. *et al.* Optical frequency comb generation from a monolithic microresonator. *Nature* **450,** 1214–1217 (2007).

a24. Huang, Y.-P. & Kumar, P. Interaction-Free Quantum Optical Fredkin Gates in Microdisks. *Selected Topics in Quantum Electronics, IEEE Journal of* **18,** 600–611 (2012).

a26. Marsh, E., John, B., Couey, J. & Wang. . ., J. Predicting surface figure in diamond turned calcium fluoride using in-process force measurement. *Journal of Vacuum . . .* <`http://link.aip.org/link/?JVTBD9/23/84/1`> (Jan. 2005).

a27. Ngoi, B. & Sreejith, P. Ductile Regime Finish Machining-A Review. *The International Journal of Advanced . . .* <`http://www.springerlink.com/index/01qynhyn3yq7vgm5.pdf`> (Jan. 2000).

a28. Blackley, W. & Scattergood, R. Ductile-regime machining model for diamond turning of brittle materials. *Precision Engineering* **13,** 95–103 (1991).

a29. Nakasuji, T., Kodera, S., Hara, S. & Matsunaga, H. Diamond turning of brittle materials for optical components. *CIRP Annals- . . .* <`http://linkinghub.elsevier.com/retrieve/pii/S0007850607610099`> (Jan. 1990).

a30. Chao, C., Ma, K., Liu, D., Bai, C. & Shy, T. Ductile behaviour in single-point diamond-turning of single-crystal silicon. *Journal of Materials Processing Technology* **127,** 187–190 (2002).

a31. Yan, J., Syoji, K., Kuriyagawa, T. & Suzuki, H. Ductile regime turning at large tool feed. *Journal of Materials Processing . . .* <`http://linkinghub.elsevier.com/retrieve/pii/S0924013601012183`> (Jan. 2002).

a32. O'Connor, B., Marsh, E. & Couey, J. On the effect of crystallographic orientation on ductile material removal in silicon. *Precision Engineering* **29,** 124–132 (2005).

a33. Yan, J., Maekawa, K., Tamaki, J. & Kuriyagawa, T. Micro grooving on single-crystal germanium for infrared Fresnel lenses. *Journal of Micromechanics and Microengineering* **15,** 1925 (2005).

a34. Fuchs, B., Syn, C. & Velsko, S. Diamond turning of lithium niobate for optical applications. *Applied Optics* **31,** 5788–5793 (1992).

a35. Yan, J., Tamaki, J., Syoji, K. & Kuriyagawa, T. Single-point diamond turning of CaF 2 for nanometric surface. *The International Journal of Advanced Manufacturing Technology* **24,** 640–646 (2004).

a36. Yan. . ., J. Crystallographic effects in micro/nanomachining of single-crystal calcium fluoride. *Journal of Vacuum Science & Technology B: . . .* <`http://link.aip.org/link/?JVTBD9/22/46/1`> (Jan. 2004).

a37. Paul, E., Evans, C., Mangamelli, A. & McGlauflin, M. Chemical aspects of tool wear in single point diamond turning. *Precision . . .* <`http://linkinghub.elsevier.com/retrieve/pii/0141635995000194`> (Jan. 1996).

a38. Maekawa, K. & Itoh, A. Friction and tool wear in nano-scale machining–a molecular dynamics approach. *Wear* **188,** 115–122 (1995).

a39. Gorodetsky, M. L. & Ilchenko, V. S. Optical microsphere resonators: optimal coupling to high-Q whispering-gallery modes. *J. Opt. Soc. Am. B* **16,** 147–154 (1999).

a41. Zhu, S., Yu, A. W., Hawley, D. & Roy, R. Frustrated total internal reflection: A demonstration and review. *Am. J. Phys.* **54** (601 1986).

a42. Court, I. N. & von Willisen, F. K. Frustrated Total Internal Reflection and Application of Its Principle to Laser Cavity Design. *Appl. Opt.* **3,** 719–726 (June 1964).

a44. Strekalov, D., Savchenkov, A., Matsko, A. & Yu, N. Efficient upconversion of subterahertz radiation in a high-Q whispering gallery resonator. *Optics letters* **34,** 713–715 (2009).

a45. Schunk, G. *et al.* Identifying modes of large whispering-gallery mode resonators from the spectrum and emission pattern. *Optics express* **22,** 30795–30806 (2014).

a47. Gopalan, V., Mitchell, T. E., Furukawa, Y. & Kitamura, K. The role of nonstoichiometry in 180° domain switching of $LiNbO_3$ crystals. *Appl. Phys. Lett.* **72,** 1981 (16 1998).

a48. Schlarb, U. & Betzler, K. Influence of the defect structure on the refractive indices of undoped and Mg-doped lithium niobate. *Phys. Rev. B* **50,** 751–757 (2 July 1994).

a49. Boyd, G. D. & Kleinman, D. A. Parametric Interaction of Focused Gaussian Light Beams. *J. Appl. Phys.* **39** (1968).

a51. Kurz, H. *et al.* Photorefractive centers in LiNbO3, studied by optical-, Mössbauer- and EPR-methods. *Applied physics* **12,** 355–368 (1977).

a52. Buse, K. Light-induced charge transport processes in photorefractive crystals II: Materials. *Applied Physics B* **64,** 391–407 (1997).

a53. Glass, A., Von der Linde, D. & Negran, T. High-voltage bulk photovoltaic effect and the photorefractive process in LiNbO3. *Applied Physics Letters* **25,** 233–235 (1974).

a54. Bryan, D. *et al.* Magnesium-doped lithium niobate for higher optical power applications. *Optical Engineering* **24,** 241138–241138 (1985).

Books

b1. Matsko, A. B. *Practical applications of microresonators in optics and photonics* ISBN: 978-1-4200-6578-7 (CRC Press, 2009).

b40. Saleh, B. E. A. & Teich, M. C. *Fundamentals of Photonics* ISBN: 0-471-83965-5 (Wiley & Sons Inc., 1991).

b43. Siegman, A. E. *Lasers* ISBN: 0-935702-11-3 (University Science Books, 1986).

Other resources

c20. McAuslan, D. L., Harris, G. I., Sheridan, E. & Bowen, W. P. *Superfluid Optomechanics: Laser Cooling Superfluid Motion* in *Integrated Photonics Research, Silicon and Nanophotonics* (2014), IM4A–1.

c25. Moore Nanotechnology Systems, L. *Nanotech 250UPL Ultra Precision Lathe* <`http://www.nanotechsys.com/machines/nanotech-250upl-ultra-precision-lathe/`> (2013).

c46. *Personal correspondence with sales representatives at United Crystals.*

c50. Risk, W. P., Gosnell, T. R. & Nurmikko, A. V. in. Chap. Supplement to Chapter 5: Derivation of the basic Boyd-Kleinman expression for SHG with circular Gaussian beams (Cambridge University Press, 2003). ISBN: 9780521521031.

Appendices

APPENDIX A

Methods

This appendix contains brief summaries of some of the experimental apparatus used for the research contained in this thesis with an emphasis on the specific models used.

A.1 Video microscopes

Many of the optics produced by the diamond turning process possess microscopic features, in the sense that they range in size from one to hundreds of microns. The substrates that these features exist on and, in the case of the smaller monolithic resonators, can be less than a millimetre across in total dimension. Accurately inspecting and positioning such devices therefore requires both precise motion control as well as high-magnification video microscopes. A number of such microscopes were used, ranging from a large Nikon stereo-microscope to a compact, budget USB microscope.

A.1.1 Nikon stereomicroscope

Inspection of machined parts is greatly aided by a high-magnification stereo-microscope. The stereoscopic vision makes the shape of small features and defects much clearer than a simple microscope. A video camera can also be mounted to record images of the machined parts.

A.1.2 Compact fixed-focal length microscope

Optical inspection of machined parts on the lathe can save considerable time by allowing defects to be determined and corrected without removing a workpiece from the lathe. A simple fixed focal length video microscope was constructed to enable visual inspection of parts between machining passes

The microscope was formed by an objective lens held a fixed distance from the CCD sensor by a lens tube that attached directly to the camera. The lens selected was an Infinistix video lens with a 20 mm working distance and 5.0X magnification, providing roughly 1 μm resolution with a 1/2” CCD sensor. The attached lens tube has a 15 mm diameter, allowing the microscope to be positioned in close proximity to other components.

The tube was then mounted onto the lathe using a kinematic base designed to allow repeatable positioning of a contact probe. This enabled the inspection camera to be easily attached and removed between cuts with repeatability of a few microns. The height was the mounting adapter was designed to locate the camera at the hight of the lathe's spindle axis and the microscope could be positioned using the lathe's motion control. This fine motion control also enables the scale of features on the image to be calibrated more precisely than by estimating the scale from the listed magnification of the lens.

A.1.3 Variable focal-length microscope

The fixed-focal length microscope is suitable for use on the lathe; however, it is cumbersome to use when precise motion control isn't available, such as when the microscope needs to be freely positioned to inspect elements on an optical bench. We therefore used a variable focal length microscope for assisting precise alignments on the optical bench.

The microscope used was an Infiniprobe S-32, with a 0.3-2.0X magnification and a working distance that varies from 28 mm to infinity. The microscope features the same camera mounting adapter as the fixed-focal length microscope and is mounted to an articulated arm that can be manually moved through six degrees of freedom via a knuckle joint, a ball joint and sliding clamps. The entire assembly is then secured to the optical bench using a moveable magnetic base. The combination of arbitrary positioning and a focal length can be varied allows easy inspection of small components on the optical bench, such as small resonators and associated coupling optics, as shown in fig. 16.6.

A.1.4 Video cameras

The primary video cameras were uEye USB CCD cameras from iDS. We used two identical cameras, each featuring a 1/2" 1280x1024 monochrome CCD sensor (Sony ICX205) with a 4.65x4.65 μm pixel size. The cameras feature a c-mount adapter which can be attached to a variety of optics.

The cameras were selected because of the availability of documentation regarding the CCD sensor as well as accessible drivers and example code. Unlike many cameras, the hardware can be fully controlled through the software user interface, with no hardware-level auto gain features. The exposure time and gain can then be held fixed so that accurate comparisons can be made between frames.

It was found that the camera drivers also worked with cheaper cameras sold by Thorlabs (likely rebranded uEye cameras), which feature the same mounting adapter. This makes for a modular system if a colour sensor is required or if the resolution needs to be upgraded at a later point.

A.1.5 Illumination

Illumination for white-light interferometry, the viewing of Newton's rings for prism alignment and general imaging of the coupling optics was provided by two fibre-coupled illuminators which could be used interchangeably. One of them is a halogen source for bright, broadband light to provide a low coherence length for white light interferometry.

The other is an RGB tuneable illuminator from scopeLED that contains three colours of LEDs at 465 nm, 521 nm, and 635 nm with linewidths of 22 nm, 32 nm, and 17 nm, respectively. The LEDs can be used independently for monochromatic illumination or simultaneously to produce an approximation of white-light. This flexibility enables control of the width of white-light interference fringes. This is useful for white-light interferometry because a longer coherence length can simplify initially finding the correct distance between the objective and a part before reducing the coherence length to measure a surface profile.

A.2 Numerical simulations

The XMDS software [1] was used extensively to model the dynamics of optical quantum memories. Some sample code for a simulation of a simple two-level gradient echo memory is provided here. The simulation solves eqs. (6.66) under the conditions of a realistic operation of the cold-atom memory. The output of the simulation is shown in fig. 9.4.

The constants and functions used in the simulation are defined in the

`<features>`

section along with options to optimise the produced code. The simulation grid

`<geometry>`

and

`<vector name>`

sections define the dimensions of the simulation and the quantities to be solved for. The

`<sequence>`

section defines the integration algorithms and the equations of motion for the system. In this case, fourth-order Runge-Kutta algorithms are being used for both dimensions, however, more complex simulations benefit from use of the adaptive fourth-fifth order Runge-Kutta algorithm for the propagation dimension. Finally, the

`<output>`

section defines how the results of the simulation should be formatted in the output file.

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- XMDS simulation for polariton interference -->

<simulation xmds-version="2">

  <name>Two_Level_GEM</name>
  <author>Geoff Campbell</author>
  <description>
    Simulation of a two-level gradient echo memory
  </description>

  <features>
    <benchmark/>
    <error_check/>
    <fftw plan="patient" />
    <auto_vectorise/>
    <globals>
      <![CDATA[
          const double pi = 3.14159;
          const double od = 2.75*63, agamma = 2*pi*0.5*6.06;
          const double delta = 2*pi*200.0, omega = 2*pi*1.705*sqrt(18.0/6.0);

          const double agammas = (pi*0.001 + agamma*omega*omega/(agamma*agamma+delta*delta));
          const complex deltas = (agamma*agamma+delta*delta)/(delta + i*agamma);

          const double input_time = 18.2, input_width = 10.2/(2*sqrt(2*log(2)));
```

```
        double gaussian( double x, double w ) { return exp( -x*x/(2*w*w) ); }
        double eta(double time) { return (time<29.0) ? 2*pi*0.143 : -2*pi*0.175; }
    ]]>
  </globals>
</features>

<geometry>
  <propagation_dimension> t </propagation_dimension>
  <transverse_dimensions>
    <dimension name="z" lattice="500"  domain="(-0.5, 0.5)" />
  </transverse_dimensions>
</geometry>

<vector name="polarisation" initial_space="z" type="complex">
  <components> S </components>
  <initialisation>
    <![CDATA[
      S = 0.0;
    ]]>
  </initialisation>
</vector>

<vector name="optical" initial_space="z" type="complex">
  <components> E </components>
</vector>

<sequence>
  <integrate algorithm="RK4" interval="60.0" steps="500" tolerance="1.0e-4">
    <samples>500</samples>
    <operators>

      <operator kind="cross_propagation" algorithm="RK4" propagation_dimension="z">
        <integration_vectors>optical</integration_vectors>
        <dependencies>polarisation</dependencies>
        <boundary_condition kind="left">
          <![CDATA[
            E= gaussian( t-input_time, input_width );
          ]]>
        </boundary_condition>
        <![CDATA[
          dE_dz = i*(od*(agamma/deltas)*E + sqrt(od)*(omega/deltas)*S);
        ]]>
      </operator>

      <integration_vectors>polarisation</integration_vectors>
        <![CDATA[
          dS_dt = -(i*eta(t)*z+agammas)*S + i*sqrt(od)*agamma*(omega/deltas)*E;
        ]]>

    </operators>
  </integrate>
</sequence>

<output format="hdf5" filename="Two_Level_GEM.dat">
  <group>
    <sampling basis="z" initial_sample="yes">
      <moments>E_real E_imag S_real S_imag</moments>
      <dependencies>optical polarisation</dependencies>
      <![CDATA[
```

```
          E_real = E.Re();
          E_imag = E.Im();
          S_real = S.Re();
          S_imag = S.Im();
        ]]>
      </sampling>
    </group>
  </output>
</simulation>
```

APPENDIX B

Material properties

This appendix contains brief summaries of the properties of various materials that are relevant to the contents of the rest of the thesis. They are restricted to the quantities needed in the studies presented here.

B.1 Lithium niobate

Lithium niobate ($LiNbO_3$) is a transparent optical crystal with electro-optical and piezo-electric properties. It is commonly used in electro-optic modulators, frequency conversion applications and optical wave-guides.

B.1.1 Crystal structure

The crystal structure gives Lithium niobate ($LiNbO_3$) its physical properties. The linear optical properties arise from the atomic density along a given axis. Its piezo-electric and electro-optic properties are the result of the positions of the Li and Nb atoms in the crystal matrix and of how they move when subject to electric fields or to strain on the crystal lattice. The dynamics of machining $LiNbO_3$ depend on the response of the structure to strain applied a various angles.

$LiNbO_3$ consists of oxygen atoms, arranged in to sheets of hexagonal lattice. Neighbouring sheets are rotated 60° with respect to each other so that the oxygens are densely packed. The interstices formed by octahedrons in the oxygen lattice are filled with one-third lithium, one-third niobium and are one third vacant. The definition of the c axis is such that the ordering in the $c+$ direction is: $\ldots$, Li, Nb, vacancy, $\ldots$.

$LiNbO_3$ is a member of the trigonal crystal system [2]; it has mirror symmetry about three planes, each separated by 60° and intersecting at an axis of three-fold rotational symmetry. The symmetry axis is referred to as the c axis and three other axes, (a_1, a_2, a_3) are chosen to be orthogonal to each of the mirror planes. The natural cleavage plane is the $\{01\bar{1}2\}$ group of planes.

By convention, crystal suppliers refer to the c symmetry axis as the z axis and employ cartesian coordinates. The x axis is chosen to be orthogonal to one of the mirror planes and the y axis is chosen to form a right-handed coordinate system. The cut axis (eg. z-cut) typically refers to the axis which is orthogonal to the face of the crystal which is polished flat. Our $LiNbO_3$, for example, is z-cut and we machine resonators in the plane orthogonal to the c axis.

Slight changes to the stoichiometry (the relative ratios of Li, Nb and O in the melt) or impurities will change the physical properties. The optical properties, in particular, can be manipulated by the addition of dopants which alter the refractive index and inhibit photorefractive effects by reducing impurities.

B.1.2 Optical properties

An alternative form of the generalised Sellmeier equation (16.9) is presented [3] to determine the refractive indices of $MgO:LiNbO_3$ for a given stoichiometric composition and MgO doping concentration:

$$n^2(\lambda, T) = 1 + \sum_j \frac{A_j}{\lambda_j(T)^2 - \lambda^2}, \tag{B.1}$$

where j denote different absorption resonances in the material.

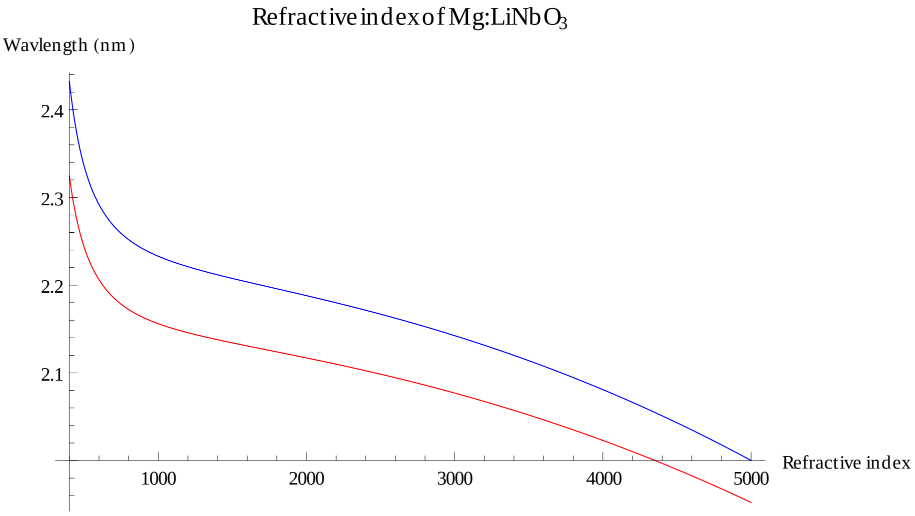

Figure B.1: The indices of refraction predicted for the ordinary (blue) and extraordinary (red) axes of $MgO:LiNbO_3$ according to the crystal composition supplied by the manufacturer.

Taking the structure of defects into account, [3] arrives at a Sellmeier equation for $MgO:LiNbO_3$ of

$$n_i^2(\lambda, T) = \frac{A_{0,i}(c_{Li}, c_{Mg})}{\lambda_{0,i}(T)^2 - \lambda^2} - A_{IR,i}\lambda^2 + A_{UV}, \tag{B.2}$$

with

$$\begin{aligned}
A_{0,i}(c_{Li}, c_{Mg}) &= \left\{ \begin{array}{ll} A_{0,i} + c_{Mg}A_{Mg,i} + (c_{thr} - c_{Mg})A_{Nb_L i,i} & : c_{Mg} < c_{thr} \\ A_{0,i} + c_{Mg}A_{Mg,i} & : c_{Mg} > c_{thr} \end{array} \right\}, \\
\lambda_{0,i}(T) &= \lambda_{0,i} + \mu_{0,i} \times (f(T) - f(T_0)), \\
f(T) &= (T + 273)^2 + 4.0238 \times 10^5 \left(\coth\left[\frac{261.6}{T + 273}\right] - 1 \right), \\
c_{thr} &= 5\% \times \frac{50\% - c_{Li}}{50\% - 48.5\%}
\end{aligned}$$

where $i = \{o, e\}$ denotes the ordinary or extraordinary axis. The first term in B.2 accounts for absorption resonances caused by the Nb sub-lattice. Due to the similarity between the resonance wavelength of the undisturbed Nb lattice and that of defect sites, only one oscillator term is needed; however, the absorption strength, $A_{0,i}(c_{Li}, c_{Mg})$, is made a function of the concentrations of Li and Mg, c_{Li} and c_{Mg} (given in molar percentage), respectively, to account for defects in the lattice. The oscillator frequency, $\lambda_{0,i}(T)$, is scaled to account for temperature with $T_0 = 24.5°\text{C}$.

The coefficients for equation B.2 that were found in [3] are shown in table B.1.2.

	n_o	n_e
$\lambda_{0,i}$	223.219	218.203
$\mu_{0,i}$	1.1082×10^{-6}	6.4047×10^{-6}
$A_{0,i}$	4.5312×10^{-5}	3.9466×10^{-5}
$A_{Nb_L i,i}$	-1.4464×10^{-8}	2.3727×10^{-7}
$A_{Mg,i}$	-7.3548×10^{-8}	7.6243×10^{-8}
$A_{IR,i}$	3.6340×10^{-8}	3.0998×10^{-8}
A_{UV}	2.6613	2.6613

Table B.1: The coefficients for the Sellmeier equation (B.2) for $MgO{:}LiNbO_3$.

The electro-optic coefficients for $LiNbO_3$, as specified by the manufacturer for a wavelength of 633 nm are summarised is table B.1.2. The unclamped coefficients include the refractive index change that is induced by the combination of the piezoelectric effect and elasto-optic effect. Similarly, the nonlinear optical coefficients for a wavelength of 1.06 μm are given.

	r_{13}	d_{31}	$r(d)_{22}$	$r(d)_{33}$	r_{51}	r_z
EO (unclamped)	10		7	33	33	18
EO (clamped)	9		3	31	28	19
Nonlinear		-5	3	-33		

Table B.2: The electro-optic and nonlinear optical coefficients for $MgO{:}LiNbO_3$ in pm/V. The electro-optic coefficients are specified for a wavelength of 633 nm, the non-linear coefficients are specified for a wavelength of 1.06 μm.

B.1.3 Optical properties

The electric permittivity tensor for $LiNbO_3$ can be expressed as

$$\epsilon_{ij} = \begin{bmatrix} \epsilon_{11} & 0 & 0 \\ 0 & \epsilon_{11} & 0 \\ 0 & 0 & \epsilon_{33} \end{bmatrix}. \tag{B.3}$$

Because the piezo-electric effect will couple the applied electric field to the strain on the crystal, the permittivity depends on whether or not the crystal is clamped to restrict its expansion. The relative electric permittivities for the clamped and unclamped cases are summarised in appendix B.1.3.

	ϵ_{11}	ϵ_{33}
EO (unclamped)	84.5	28.9
EO (clamped)	44.7	27.0

Table B.3: The relative electric permittivity for $LiNbO_3$. The values are averaged from the sources summarised in [2]

B.2 Fluoride crystals

Calcium fluoride (CaF_2) and magnesium fluoride (MgF_2) are highly transparent optical crystals that are used in applications that require very high transmission. The relatively low refractive index and low dispersion of the crystals means that they are often used was uncoated optical windows. Fluoride crystals have a transmission spectrum into the UV and have found applications in semi-conductor lithography, where short wavelengths are used to achieve the finest resolution possible.

B.2.1 Calcium fluoride

The refractive index of CaF_2 can be calculated from the Sellmeier dispersion formula

$$n^2 - 1 = \sum_{i=1}^{4} \frac{A_i \lambda^2}{\lambda^2 - B_i} \tag{B.4}$$

with the coefficients given by appendix B.2.1 [4]. At the wavelengths of interest for this thesis, the refractive indices are 1.4354 at 532 nm and 1.4285 at 1064 nm. For high purity samples, the bulk photon absorption can be very low. The longest photon lifetime measured in CaF_2 is 1 ms [5] and was possibly limited by surface scattering.

A_1	$4.430595147 \times 10^{-1}$	B_1	$1.733873966 \times 10^{-3}$
A_2	$4.454624348 \times 10^{-1}$	B_2	$7.939087382 \times 10^{-3}$
A_3	$1.502595301 \times 10^{-1}$	B_3	$1.234337898 \times 10^{-2}$
A_4	$8.859807728 \times 10^{0}$	B_4	$2.751117861 \times 10^{-3}$

Table B.4: The Sellmeier coefficients for CaF_2 at 20°C

B.2.2 Magnesium fluoride

MgF_2 is a similar transparent optical crystal to CaF_2 ; however, it is birefringent. The refractive indices [6] for the wavelengths used in this thesis are summarised in appendix B.2.2

	532 nm	1064nm
n_e	1.3908	1.3848
n_o	1.3789	1.3732

Table B.5: Some refractive indices for MgF_2 at 19°C. From the Sellmeier equations in [6].

Glass	Sellmeier equation ($n^2 - 1$)	n(532 nm)	n(1064 nm)
N-BK7	$\frac{1.0396\lambda^2}{\lambda^2-0.0047} + \frac{0.4079\lambda^2}{\lambda^2-0.0135} + \frac{0.8975\lambda^2}{\lambda^2-97.9340}$	1.5195	1.5066
UV fused silica	$\frac{0.6962\lambda^2}{\lambda^2-0.0132} + \frac{0.3137\lambda^2}{\lambda^2-0.0623} + \frac{1.8988\lambda^2}{\lambda^2-155.2363}$	1.4607	1.4450
SF11	$\frac{1.7376\lambda^2}{\lambda^2-0.0132} + \frac{0.3137\lambda^2}{\lambda^2-0.0623} + \frac{1.8988\lambda^2}{\lambda^2-155.2363}$	1.7948	1.7539

Table B.6: The properties of the optical glasses used as coupling prisms.

	A	B	C (10^{-2})	D	F
n_o	1.73358749	0.96464345	1.94325203	1.82831454	120
n_e	1.35859695	0.82427830	1.06689543	0.14429128	120

Table B.7: The coefficients for the dispersion of CaF_2 , as described by eq. B.5

B.3 Prisms

B.3.1 Optical glasses

The three optical glasses used for coupling prisms in various experiments are N-BK7, UV fused silica, and SF11. Table B.3.1 lists the refractive indices of these glasses. The equations are taken from the Thorlabs datasheets for the materials, as this is the supplier of the prisms that we used.

B.3.2 Calcite

The refractive indices for the ordinary and extraordinary rays in CaF_2 are provided in [7] by the dispersion equation

$$n^2 = A + \frac{B\lambda^2}{\lambda^2 - C} + \frac{D\lambda^2}{\lambda^2 - F}, \tag{B.5}$$

with coefficients listed in table B.3.2. The refractive indices of calcite for the wavelengths relevant to second harmonic generation of 1064 nm light are presented in appendix B.3.2.

B.3.3 Rutile

Rutile is a hard optical crystal with an exceptionally high index of refraction, given by the Sellmeier equation [8]:

$$n_e^2 = 7.197 + \frac{0.3322}{\lambda^2 - 0.0843} \tag{B.6}$$

$$n_o^2 = 5.913 + \frac{0.2441}{\lambda^2 - 0.0803}. \tag{B.7}$$

The refractive indices for 532 nm and 1064 nm light are listed in table B.3.3. Rutile absorbs

	532 nm	1064nm
n_e	1.4882	1.4796
n_o	1.6628	1.6425

Table B.8: The refractive indices for calcite at 532 nm and 1064 nm.

	532 nm	1064nm
n_e	2.9780	2.7412
n_o	2.6678	2.4789

Table B.9: The refractive indices for rutile at the relevant wavelengths for second harmonic generation of 1064 nm light.

strongly in the short wavelengths of the visible spectrum, giving it a yellow appearance and therefore results in some loss for 532 nm light even over the short optical path length of the coupling prism.

APPENDIX C

The Moore Nanotech 250UPL

The diamond turning machine used for the work in this thesis is the Moore Nanotech UPL250, ultra-precision lathe [9]. While the technology is commercial rather than developed in-house, a brief overview is warranted due to its central role in the work of my candidacy.

C.1 The axes of motion

The UPL250 is a collection of high-performance motion-control components assembled on a vibration isolation system. The model that we use has two linear axes (labeled X and Z), a rotary axis mounted normal to the X-Z plane on the X axis and a work spindle mounted on the Z axis with its axis oriented parallel to the X axis. The axes are mounted on a granite table which absorbs acoustic vibrations and is raised on pneumatic cushions to limit the transmission of vibration from the floor on which the lathe sits.

C.1.1 The linear axes

The motion control used for shaping a workpiece is primarily performed by the two linear axes of the lathe. These consist of linear hydrostatic box-way slide bearings. High-pressure oil from a hydraulic pump is forced between rectangular guideways which are fixed to the machine base and the slides that carry the axes. The gap between the slides and guideways is only a few microns across which provides a high stiffness with very low friction. The guideways are specified to have less than 200 nm of deviation from a straight line over the full 200 mm of travel.

Position feedback is provided by diffraction gratings etched into glass reference plates. The analog signals generated by the holographic gratings are digitally encoded and provide a resolution of 34 picometers. Linear magnetic servo motors are used to maintain precise motion control of the axes by closed-loop feedback via a tuned, digital PID controller.

C.1.2 The work spindle

The workpiece is held on an air-bearing spindle which is equipped with liquid cooling to maintain a narrow temperature range to avoid motion due to thermal expansion during operation. The air bearing, using compressed air to maintain a gap of a few microns between the spindle shaft and casing, provides an axial stiffness of 228 N/μm and a radial stiffness of 98 N/μm. The position accuracy is less than 12 nm in both the axial and radial direction during operation.

The spindle drive motors operate in a closed-loop mode and can rotate the spindle at up to 10 000 rpm or provide precise slow positioning with an accuracy of 1 arcsecond. The ability to use the spindle servo motors in a precision motion control mode allows the machining of non-cylindrically symmetric pieces by synchronising the spindle position with the motion of the linear axes.

The workpiece is held on the spindle by a vacuum chuck. The pressure of the vacuum can be set to reach a suitable balance between holding force and sag in the workpiece. In the usual case of using an additional holder between the chuck and the piece to be machined the maximum vacuum pressure is used.

C.1.3 The rotary axis

The rotary (B) axis is used to rotate the diamond tool during cutting or to swap out different tools. It uses a vertically oriented air-bearing axis with stiffnesses equal to those of the work spindle. The positioning accuracy is 1 arcsecond. Motion of the B axis can be synchronised with that of all the other axes.

C.1.4 The motion controller

Motion control is provided by a Delta Tau PMAC motion controller. The controller is provided with digitised encoder signals from each of the lathe's axis. The controller performs a digital calculation of the PID feedback required to maintain control of the position of each axis through a programmed move. Movements are synchronised across all axes to produce a programmed tool path. Various types of interpolation between points, linear, circular or spline, can be computed by the controller.

The output of the controller is fed into the drive amplifiers and subsequently to the servo motors.

C.2 Non-cylindrically symmetric parts

On of the key features of the ultra-precision lathe is to synchronise movements of the linear axes to the position of the work spindle during cutting. This allows non-cylindrically symmetric components to be produced on the lathe. A tool path is specified as a list of $\{C, X, Z\}$ coordinates for the position of the three axes involved in the cut, the C axis being the angle of the work spindle in degrees.

While the software provided with the lathe is capable of producing a set of commonly used shapes, we opted to write the tool paths using external software. For cutting on the face of a part, this is done by programming a spiral path using $\{C, X\}$ polar coordinates and mapping each point on the path to a height in Z. For machining on the outer diameter, the same is done using $\{C, Z\}$ cylindrical coordinates and map that provides the height in X.

The speed of the cut is programmed by setting the time between points. In general, the time is short enough that the lathe runs at the maximum speed allowed by acceleration of the axes. In this case, the motion control is not as precise as the resolution of the encoders but is the convolution of the programmed path with some transfer function which is determined by the maximum allowable deviation from the programmed path which the controller decides is acceptable. Fine control over these parameters is possible by altering parameters set in the motion control board, but in practice, it is simpler to change the density of sampling points on the programmed tool path until an acceptable form accuracy is obtained.

APPENDIX D

Other machining projects

In addition to fabricating $LiNbO_3$ monolithic resonators for investigating non-linear frequency, I worked on a number of other projects that leverage the capabilities of the ultra-precision lathe. These projects ranged in scope from producing simple machined optics for collaborating research groups to performing substantial original research. The original work was directed towards producing high-quality spiral phase plates for generating high-charge optical vortices. This project was successful in generating enough interest that we proceeded to produce multiple spiral phase plates for collaborators working in diverse fields. Other work done for collaborating research groups involved producing other novel diffractive optics, whispering gallery resonators, and a silicon flexure mount for opto-mechanical investigations. Here, I present a brief summary of some of these projects.

D.1 Spiral phase plates

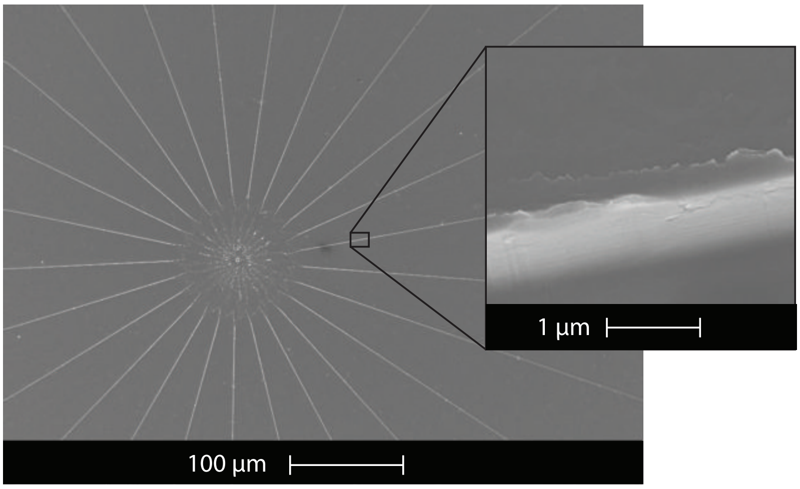

Figure D.1: SEM image of a spiral phase plate. The defect in the center is caused by the radius of the diamond as it approaches the phase singularity at the center of the optic.

Optical vortices are light beams that carry orbital angular momentum and are characterised by a phase that increases azimuthally about a singularity at the centre of the beam[10, 11]. These beams have a variety of possible applications, including include opti-

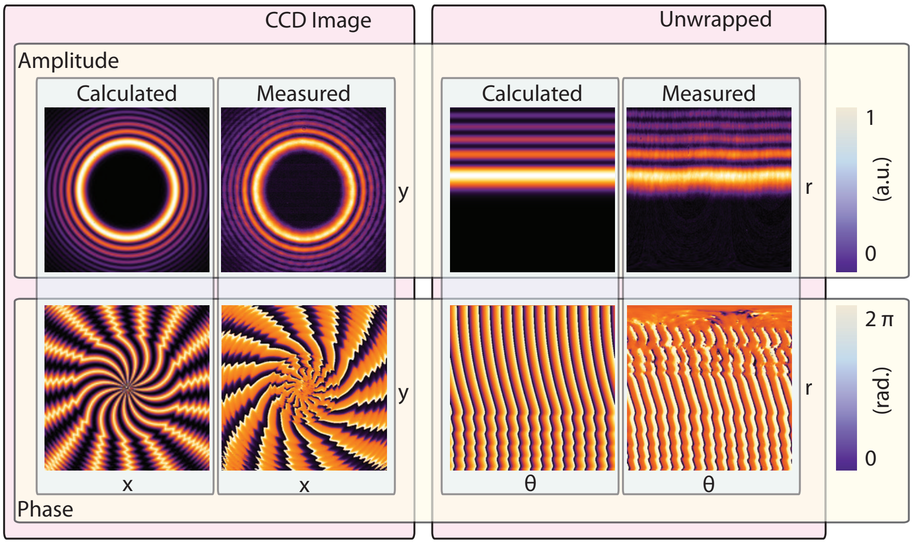

Figure D.2: Phase and amplitude measurements of a charge 15 optical vortex along with calculated amplitudes for an ideal vortex.

cal trapping and manipulation of colloidal particles[12] and cold atoms[13], the generation of Bose-Einstein condensate vortices[14], phase-contrast microscopy[15], image filtering for extra-solar planet detection[16], stimulated emission depletion (STED) microscopy[17] and the encoding of optical quantum information[18].

One method to apply the required phase shift is to use a computer-controlled spatial-light modulator (SLM)[19]. Another approach is to use lithographic processes to construct a phase mask[20–22]. We demonstrate that diamond-turned phase mirrors can produce high-quality optical vortices with topological charges that far-exceed those that can be achieved with more conventional methods. The spiral phase mirrors, also referred to as phase plates, impart the azimuthal phase onto a beam that is reflected off of them at normal incidence by featuring a helical pattern on the surface.

The phase-plates are produced by directly cutting a helical ramp onto an aluminium substrate. A form accuracy of well under an optical wavelength is achieved by dividing the plate into many segments, each applying a fraction of the total angular momentum. The steps are separated by a sharp step with a height that is an integer number of wavelengths. An electron-microscope image of the resulting structure is shown in fig. D.1.

D.1.1 High-quality vortices

We demonstrated that the produced optical vortices have an extremely high fidelity with ideal vortices by performing interferometric measurements on them. The amplitude of the vortices was recorded on a CCD camera, as were interference patterns formed with a plane reference field. This allowed a reconstruction of the complex amplitude of the beam. Figure D.2 shows the results. The agreement between the calculated ideal vortices and measured vortices is excellent. In particular, the very dark region in the centre of the vortices indicates that there is very little undiffracted light or scattering. the results are published in [23].

In addition to our own research, we supplied a number of vortex phase mirrors to other research groups. These included low-charge and half-charge vortex mirrors for research into optical phase singularities at the Australian National University and low-charge vortex

Figure D.3: Image of a charge 10,000 vortex mirror. The optic consists of 125 steps, each imparting a phase shift of $80 \times 2\pi$.

mirrors for hollow dipole trapping beams used to load rubidium into hollow fibres in the group of Prof. Andre Luiten.

D.1.2 High-charge vortices

The method of fabrication for the vortex plates is not limited in term of the height of the helical steps, and therefore the charge of the vortex. We investigated producing ultra-high-charge vortices by increasing the step height. Figure D.3 shows a reflection off the surface of a charge 10,000 vortex plate. The details of the characterisation of the high-charge phase plates in published in [24]. The project has also resulted in a collaboration with the research group of Prof. Anton Zeilinger to investigate the entanglement of photons carrying large angular momentum.

D.2 Diffractive optics

Similarly to the spiral phase mirrors, arbitrarily-shaped diffractive optics can be produced.

D.2.1 Mode-converter

We produced a number of diffractive optics for a company, CAILabs, that is working on spatial mode-converters for use in the telecom industry. The mode converters are designed to take multiple nearby Gaussian modes and combine them into a complex pattern that can be sent through a multi-mode optical fibre. The mode-shaping can be reversed at the far end of the fibre to restore the original separated beams. If successful, the project would multiply the information capacity of optical fibre networks.

The mode-converters operate by reflecting the input multiple times on diffractive optics with complex height profiles. The diffractive optics impart a new phase profile to the beam, which then propagates through a lens. The lens images the spatial Furier transform of the beam onto the next optic. Through a series of Fourier transforms and phase masks an arbitrary spatial transformation can be applied to the beam. We cut a diffractive

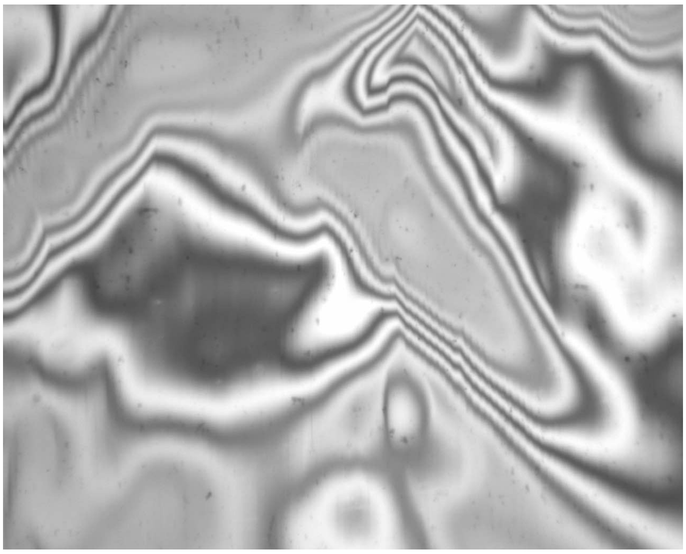

Figure D.4: Phase and amplitude measurements of a charge 15 optical vortex along with calculated amplitudes for an ideal vortex.

that includes six regions that will be illuminated to complete the mode transformation, corresponding to six iterations of a spatial Fourier transform. Figure D.4 shows a white-light interferometer measurement of one of these regions showing a complex phase profile. The results are published in [25].

D.2.2 Engineered scattering centre

The method of separating the probe and control fields in the cold-atom memory experiment discussed in section 7.4.4 was not the only method tried. Because the control field is collimated with a large beam waist as it passes through the memory, it is focused by the lens that is used to collimate the divergent probe field after the memory. At the focus, the control field can therefore be eliminated by placing a scattering centre at the focal point. A small conical mirror was machined onto a large aluminium optical flat to accomplish this. The reflective cone ensures that the control field is scattered well away from the collection optics for the probe field. The method was proven to be effective for filtering nearly all of the control field, however, not as effective as using a small angle to separate the beams, which is what was ultimately used. An white-light interformeter image of the cone is shown in fig. D.5.

D.3 Whispering gallery resonators

The production of crystalline whispering gallery resonators was one of the primary goals of research involving the ultra-precision lathe and one that was met quite successfully. Once we developed the technical expertise to produce high optical quality factor resonators, we were able to supply them to other research groups for use in at least three avenues of investigation.

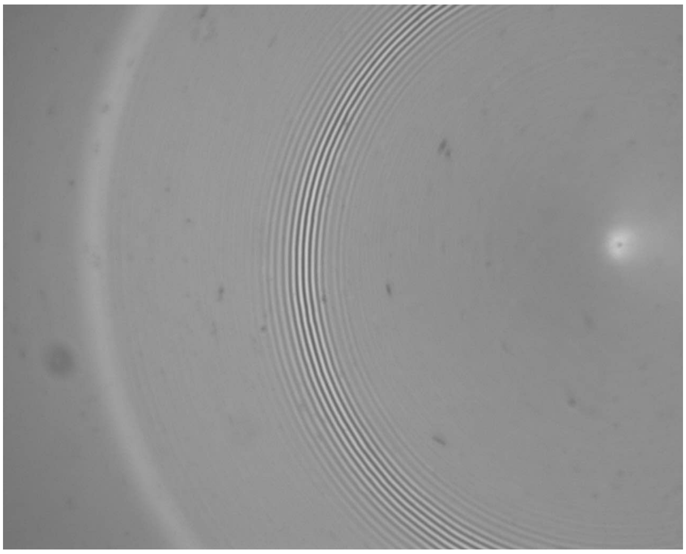

Figure D.5: A white-light interferogram of a 500 μm cone on an optical flat. This optic was used as a scattering centre to remove a focused control field from the signal field in cold-atom memory experiments.

D.3.1 High-Q Mg:F_2 resonators for temperature sensing

CaF_2 and Mg:F_2 have extremely low bulk absorption making them well-suited to sensing applications. We supplied resonators of both materials to the group of Prof. Andre Luiten at the University of Adelaide to work on using them for sensing applications. The resonators were of the same geometry as those discussed in section 15.4. The collaboration resulted in an optical sensor capable of resolving nano-Kelvin scale temperature variations. The principle is based on the differential change in the refractive index of Mg:F_2 for light at different frequencies. The resonator was probed with 1064 nm and 532 nm light and the change in the separation of resonances for the two frequencies was used to infer the temperature of the material in the optical mode. Feedback can then be applied to the mode temperature by modulating the optical input power, stabilising the temperature to below thermodynamic limits. The technique has applications in laser frequency stabilisation and the results are presented in [26].

D.3.2 Chalcogenide resonator for comb generation

The non-linearity present in optical crystals can be used perform frequency conversion such as four-wave-mixing (4WM). In this process, two photons from a bright field are converted to two photons of different frequencies whose energy adds up to the energy of the two original photons. For weak non-linearities, large intensities are required to produce significant non-linear conversion. Whispering gallery resonators provide a chance to substantially amplify the optical intensity of a resonant optical mode thanks to the large finesse that is achievable. In such resonators, two photons in a resonant optical mode can undergo a 4WM process to create a photon in each of the neighbouring modes because the frequency spacing is equidistant so energy is conserved. The large finesse associated with all four photons involved strongly amplifies the process. When the 4WM is sufficient to convert a large fraction of the pump light, the process can cascade to form a frequency comb [27].

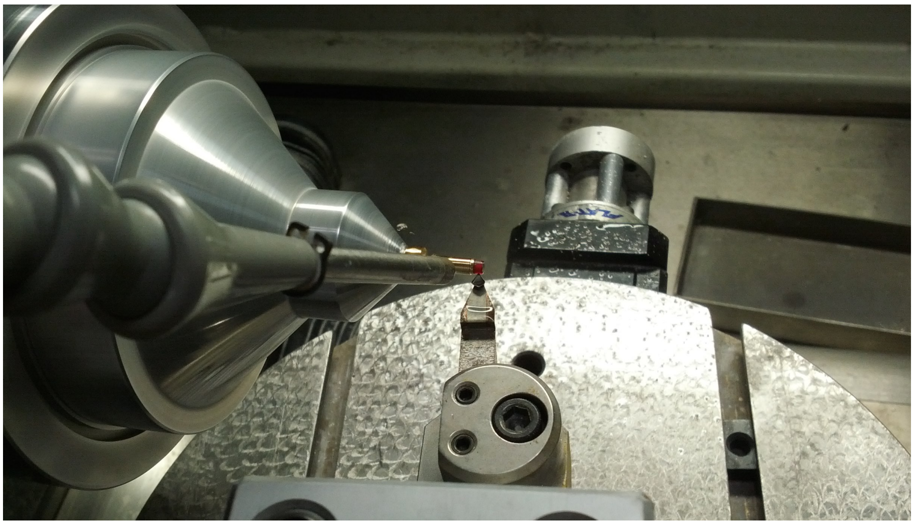

Figure D.6: Machining a chalcogenide whispering gallery resonator.

Comb generation using CaF_2 whispering gallery resonators has been successfully used to generate octave-spanning combs in the visible spectrum and the near-infra-red [27]. We produced a chalcogenide glass whispering gallery resonator for the research group of Prof. Barry Luther-Davies at the Australian National University to investigate comb generation in the far infra-red spectrum. The large non-linearities present in chalcogenide may afford a chance to produce a mode-locked frequency comb provided that the dispersion can be engineered. Such dispersion engineering can be accomplished by tailoring the mode structure of the resonator, and we supplied a resonator that was designed by Prof. Luther-Davies' team to support only a single transverse spatial mode in the frequency spectrum of interest. Figure D.6 shows a picture of chalcogenide being machined.

D.4 Other parts

D.4.1 Hemispherical mirror

As part of a collaboration with the research group of Rainer Blatt in Innsbruck, we have been working to fabricate a high-precision reflective concave hemisphere for use in ion trapping experiments. The project is to trap an ion at the center of the hemisphere and use the reflection from the hemisphere to modify the spontaneous emission rate from the ion. Using the rotary axis of the lathe to define an arc-shaped cut, we were able to machine concave hemispheres to better than wavelength form accuracy.

D.4.2 Silicon flexure

The observation of opto-mechanical effects typically requires a high mechanical quality factor for a mechanical resonance so that weak forcing from optical radiation pressure can be resolved. One of the limiting factors for mechanical quality factors is loss due to the intrinsic properties of the material that a flexure is made from. To reduce such losses, stiff crystalline materials can be used as they exhibit very low intrinsic loss damping.

We machined a high-quality factor flexure made from silicon for opto-mechanical research at the Centre for Gravitational Physics at the Australian National Univeristy. The

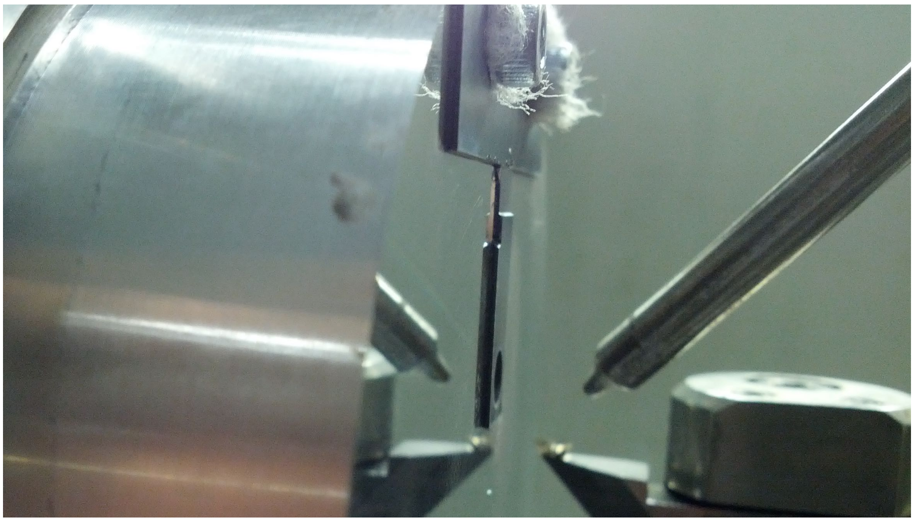

Figure D.7: Machining a silicon flexure. Care was taken to avoid fracturing the thin substrate when attaching it to the mounting flat.

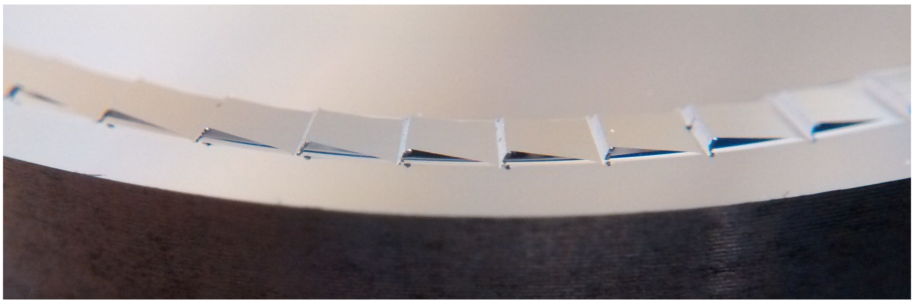

Figure D.8: A stepped variable delay line. The optic has a similar shape to the vortex mirrors but with a much larger depth.

flexure is a T-shaped piece of pure single-crystal silicon designed to be mounted by the horizontal section such that the vertical section can felx freely. We were supplied with silicon T-shapes that ranged from 300 μm to 500 μm thick. To form the flexure, we reduced the thickness of the flexure at the location where the vertical section meets the horizontal mounting. The thickness was reduced to 150 μm, taking care to use nearly ductile material removal so that stresses on the thin flexure were kept to a minimum. Figure D.7 shows the machining process.

D.4.3 Variable delay ring

Similar to the high-charge vortex mirrors, we produced a circular stepped ramp on an aluminium substrate. The optic will be rapidly rotated for use as a variable delay-line in fibre sensing applications. Each step is 250 μm in height, far higher than the steps in a vortex mirror, and forms a 10° angle relative to the flat surface of the substrate. The 22 mm diameter optic is shown in fig. D.8

Bibliography - Appendices

1. Dennis, G. R., Hope, J. J. & Johnsson, M. T. XMDS2: Fast, scalable simulation of coupled stochastic partial differential equations. *Computer Physics Communications* **184,** 201–208. ISSN: 0010-4655 (2013).
2. Weis, R. & Gaylord, T. Lithium niobate: summary of physical properties and crystal structure. *Applied Physics A: Materials Science & Processing* **37,** 191–203 (1985).
3. Schlarb, U. & Betzler, K. Influence of the defect structure on the refractive indices of undoped and Mg-doped lithium niobate. *Phys. Rev. B* **50,** 751–757 (2 July 1994).
4. Incorporated, C. *Calcium Fluoride Product Information Sheet - Code 9575* <`http://www.corning.com/WorkArea/linkit.aspx?LinkIdentifier=id&ItemID=34271`> (2014).
5. Grudinin, I., Matsko, A. & Savchenkov, A. Ultra high Q crystalline microcavities. *Optics.* <`http://linkinghub.elsevier.com/retrieve/pii/S0030401806002136`> (Jan. 2006).
6. Dodge, M. J. Refractive properties of magnesium fluoride. *Appl. Opt.* **23,** 1980–1985 (June 1984).
7. Ghosh, G. Dispersion-equation coefficients for the refractive index and birefringence of calcite and quartz crystals. *Optics Communications* **163,** 95–102. ISSN: 0030-4018 (1999).
8. DEVORE, J. R. Refractive Indices of Rutile and Sphalerite. *J. Opt. Soc. Am.* **41,** 416–417 (June 1951).
9. Moore Nanotechnology Systems, L. *Nanotech 250UPL Ultra Precision Lathe* <`http://www.nanotechsys.com/machines/nanotech-250upl-ultra-precision-lathe/`> (2013).
10. P. Coullet, L. Gil & F. Rocca. *Opt. Commun.* **73** (403 1989).
11. S. Franke-Arnold, L. Allen & M. Padgett. *Laser & Photon. Rev.* **2** (299 2008).
12. D. Grier. *Nature* **424** (2003).
13. T. Kuga, e. a. *Phys. Rev. Lett.* **78** (1997).
14. M. F. Andersen, e. a. *Phys. Rev. Lett.* **97** (2006).
15. S. Bernet, A. Jesacher, S. Fürhapter, C. Maurer & M. Ritsch-Marte. *Opt. Express* **14** (3792 2006).
16. Lee, J. H., G. Foo, Johnson, E. G. & Swartzlander, G. A. *Phys. Rev. Lett.* **97** (2006).

17. Willig, K. I., Rizzoli, S. O., V. Westphal, R. Jahn & Hell, S. W. *Nature* **440** (440 2006).
18. A. Mair, A. Vaziri, G. Weihs & A. Zeilinger. *Nature* **412** (313 2001).
19. S. Sundbeck & I. Gruzberg. *Opt. Lett* **30** (447 2005).
20. Kennedy, S. A., Szabo, M. J., H. Teslow, Porterfield, J. Z. & Abraham, E. R. *Phys. Rev. A* **66** (043801 2002).
21. K. J. Moh, e. a. *App. Opt.* **45** (1153 2006).
22. Kotlyar, V. V., Kovalev, A. A., Skidanov, R. V., Khonina, S. N. & J. Turunen. *App. Opt.* **47** (6124 2008).
23. Campbell, G., Hage, B., Buchler, B. & Lam, P. K. Generation of high-order optical vortices using directly machined spiral phase mirrors. *Appl. Opt.* **51,** 873–876 (Mar. 2012).
24. Shen, Y. *et al.* Generation and interferometric analysis of high charge optical vortices. *Journal of Optics* **15,** 044005 (2013).
25. Genevaux, P. *et al. 3 Modes transmission using hybrid separation with high mode selectivity and low losses spatial mode multiplexer* in (IEEE, 2014), 1–3. doi:`10.1109/ECOC.2014.6964119`.
26. Weng, W. *et al.* Nano-Kelvin Thermometry and Temperature Control: Beyond the Thermal Noise Limit. *Phys. Rev. Lett.* **112,** 160801 (16 Apr. 2014).
27. Kippenberg, T. J., Holzwarth, R. & Diddams, S. Microresonator-based optical frequency combs. *Science* **332,** 555–559 (2011).