Atom-light interfaces for quantum information processing

Jesse L. Everett

A thesis submitted for the degree of Doctorate of Philosophy in Physics of The Australian National University

March 2018

Declaration

This thesis is an account of research undertaken between January 2013 and January 2018 at The Department of Quantum Science, Research School., 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.

Jesse L. Everett
2018

Acknowledgements

This thesis would not exist without the work of many people. The research presented here was made possible firstly by my supervisors, Ben Buchler and Ping Koy Lam. You organised an amazing group of people and an impressive research program in quantum atom-optics.

I started my PhD towards the end of a dedicated effort to build a magneto-optic trap suitable for running a quantum memory. The work done by many on this setup made possible the experiments in cold atomic ensembles presented here.

I explored several topics during my PhD, but I would especially like to thank the people that worked with me on stationary light: Geoff, Young-Wook, Pierre, and Daniel. Your contributions to the theoretical and experimental work made Raman stationary light what it is.

Thank you to everyone who made the practical parts of the research possible: Amanda Haines and Lynne Christians, our group administrators, for organising travel, equipment, people, Tim Tams, and coffee. Thanks also to Neil Hinchey, Neil Devlin, and Paul McNamara for building our equipment in the electronic and mechanical workshops.

I would like to thank everyone I worked with in the lab: Mahdi Hosseini, Olivier Pinel, Julien Bernu, Jiao Geng, MingTao Cao, Daniel Higginbottom, Pierre Vernaz-Gris, Geoff Campbell, Young-Wook Cho, Su Jian, Anthony Leung, Aaron Tranter, and Karun Paul. Your enthusiasm and camaraderie made the experimental work enjoyable.

Merci á Pierre pour notre premier essai de lumière stationnaire, et merci aussi á Julien et LKB pour votre hospitalité.

Thank you to those that read drafts of this thesis and offered useful feedback, and best of luck to the continuing PhD students of the group.

Finally, thank you Mum and Dad for your support throughout my education and the start of my career as a scientist. Thank you Tanja for looking after me so well during my PhD.

This thesis is based on research I conducted while receiving an Australian Government Research Training Program Scholarship. My research was also supported by the Australian Research Council Centres of Excellence program through the Centre for Quantum Computation and Communication Technology (CE110001027).

Abstract

The emergence of quantum physics from the page to the lab and the world at large is an exciting development of recent years. The prospects of absolutely secure communication and efficient simulation of physical systems have spurred great human effort into understanding these possibilities and turning them into realities.

Photons are the most easily manipulated quantum particles and are a promising candidate for implementing these technologies. Limitations of photons include the difficulty of keeping objects that move at the speed of light, and producing strong interactions between particles that do not normally interact. The work presented in this thesis is motivated by the possibility of overcoming these limitations.

The ability to faithfully store and reproduce a quantum state is essential for many quantum information technologies. Quantum memories for light have been developed over the last two decades to provide this ability. The group at the Australian National University developed the gradient echo memory (GEM): A quantum state of light can be controllably stored and released from an atomic ensemble by the use of additional optical fields and magnetic field gradients. This scheme was previously shown to preserve the quantum characteristics of the light.

We used the GEM scheme with a cold rubidium ensemble to create the first optical memory that simultaneously beat the no-cloning limit, a benchmark for many of the technologies relying on quantum memories, and the loss rate for a delay line composed of optical fibre.

We also created an analogue to a pulsed optical resonator using GEM with a warm rubidium vapour. This was done by replacing the circulating optical field of a resonator with light stored in the memory, and replacing the coupling of light into and out of that circulating mode with storage and recall from the memory. The bandwidth and repetition rate of this resonator were rapidly tunable as they were controlled by external optical and magnetic fields.

We worked on implementing GEM with strings of thousands of atoms strongly coupled to the evanescent field of an optical nanofibre. This raised new possibilities for creating a true random access memory that would allow a more flexible use of the multi-mode capacity of GEM.

We developed the theory for a novel type of stationary light in the gradient echo memory. Our stationary light scheme relies on the destructive interference of counter-propagating optical fields throughout the memory. The optical intensity scales with optical depth, as with other forms of stationary light. However, as the destructive interference could be set up over a much greater distance, more of the optical depth is available for generating stationary light.

Finally, we studied how a control-phase gate for single-photon optical states could be implemented using a nonlinear interaction with stationary light. The stationary light generated by one state modulates the phase of another state stored in the memory. The second state modifies the stationary light, also producing a back-action on the first state and generating the required cross-phase shift.

Contents

List of Figures

Chapter 1

Introduction

> Sir Isaac Newton first dissected a ray of light, and proved that it was composed of several colors, but the subject has recently been elevated into a special science, called *spectral analysis*, by the splendid discoveries of the two German professors, Kirchoff and Bunsen. Professor Kirchoff used four prisms of very perfect workmanship to examine the solar spectrum through a telescope having a magnifying power of 40. He saw whole series of nebulous bands and dark lines, and a new field of vision, like that first developed by the microscope, was opened up. These dark lines it is conjectured have been made to reveal the chemical composition of the sun's atmosphere.
>
> *Scientific American*, 1863

The use of light to study the nature of matter helped propel the modern scientific revolution. These investigations also revealed the first hints of the quantum nature of light, though the understanding of light being composed of photons was ultimately prompted by the observation of other light-matter interactions. Quantised light explained the overlying spectrum of black-body radiation in thermal spectra and the wavelength cut-off of the photoelectric effect.

In the 1920s, quantum mechanics was developed as a consistent explanation of the behaviour of quantum systems. Over the intervening decades this was developed into a nearly complete description of the universe. The realisation that only a machine exploiting the weird behaviour of quantum mechanics could efficiently model real systems helped inspire modern understanding of the relationship between quantum information and reality. Quantum information theory has many implications for fundamental physics, but has also inspired quantum communication and computing as applications.

Quantum communication is a focus of work in quantum atom optics. Light is the preferred carrier of information in global communication infrastructure because it is travels notoriously quickly and is easy to create. Light will also be the preferred carrier for quantum communication. Sharing entanglement between distant locations will allow for absolutely secure communication, but will require the use of optical quantum memories at repeater stations to overcome the loss of light in optical fibre networks.

The non-classical distributions that arise from entanglement and superposition are expected to allow exponential speed increases for some computations. This includes the efficient simulation of physical systems, opening up new fields in material and biological sciences. From the perspective of fundamental physics, the idea of using controllable quantum systems to explore the properties of matter could be pursued at such scales and accuracy that would allow an explanation of how the classical world emerges from the quantum mechanical world, and perhaps answer outstanding questions about the relationship between quantum mechanics and general relativity.

The "nebulous bands and dark lines" in the spectrum of sunlight are due to atoms in the sun's atmosphere absorbing very particular wavelengths of light. This narrow bandwidth makes atoms an excellent platform for interacting with light as the interactions can be tuned very precisely. The identical energy levels of each atom are a further bonus as the light interacts similarly with all the atoms. These interactions can then be strengthened by using a large collection of atoms; an atomic ensemble.

We first study the use of atoms as a quantum memory for light. An optical quantum memory preserves the complete quantum state of the light, even down to its entanglement with other objects. Quantum memories are an integral part of quantum information processing, and linking up systems in a quantum internet will require the encoding of quantum information in optical states. Quantum communication will require the use of quantum repeaters, which allow the distribution of entanglement over distances much longer than a single entangled pair of photons can be reliably sent. Repeaters act as waypoints, allowing entanglement to be created between distant nodes by combining shorter stages, a process that requires reliable memories. The gradient echo memory (GEM) is an optical quantum memory that has been developed primarily at the Australian National University (ANU). In Part I of this thesis I present some extensions to the memory scheme, as well as a demonstration of GEM as a quantum memory in cold atoms. This implementation of GEM currently holds the record for efficiency as an optical quantum memory.

A property of light that makes it useful for communication is that it does not interact with light. This property also makes light less useful for quantum computing; many quantum computing schemes require nonlinear interactions between quantum states. Light does interact due to virtual particles in the vacuum, but real atoms act as a better intermediary. It is possible to use the interaction of a single photon with an ensemble of atoms to measurably change the interaction of a second photon with the same atoms and so entangle the two photons. The current challenge is to make this interaction more reliable and reduce the chance of losing a photon.

This is the motivation for our next area of study: stationary light. Stationary light is a type of light-matter interaction based on electromagnetically-induced transparency (EIT) that completely traps the light within the matter. The light is not completely converted to excitations of the matter, but instead some portion of it exists as an optical field. Over time the trapped field can apply an enhanced interaction compared to freely propagating light. In Part II of this thesis I present a novel form of stationary light, Raman stationary light, based on the off-resonant interaction of light with atoms. The Raman stationary light is qualitatively different, possibly allowing greater nonlinearities than other forms of stationary light. The atom-light interaction under the stationary light conditions also allows for a type of memory that could not normally be implemented without a cavity. The behaviour of the stationary light and the memory are studied both theoretically and experimentally.

When atomic ensembles and multiple classical and quantum light fields are combined, there is a much greater potential for complex interactions. We are motivated to study nonlinear interactions for optical quantum computing, so we focus on the simpler interactions. In Part III of this thesis I analyse how the Raman stationary light will behave when used for cross-phase modulation. In this scheme, the state generating stationary light is written such that it will interact with another stored state. The nonlinearity occurs by the stationary light shifting the energy levels of atoms on which the second state is stored, imparting a phase shift of the state. The stationary light is in turn affected by

back action from the target state. The stationary light experiences dispersion as it crosses, limiting the possible nonlinearity but also generating the desired cross-phase modulation. This is a simple illustration of a general limitation on cross phase modulation that ultimately derives from the causal, local nature of the interaction.

In summary, this thesis contains three perspectives on atomic ensembles as interfaces for quantum information processing. Part I contains my work with a quantum memory for light, the gradient echo memory. Part II contains my work on theoretical derivation and experimental demonstration of Raman stationary light, along with an optical memory based on the same interaction. Part III contains my theoretical work exploring the use of Raman stationary light to enhance nonlinear interactions. Each part contains a literature review to provide context, and the relevant background theory to allow a complete understanding of the work.

Chapter 2

Quantum theory for atom-light interaction

> Only chance can speak to us.
>
> Milan Kundera, *The Unbearable Lightness of Being*

Light and atoms are quantum objects and an accurate description of their behaviour can only be derived in the framework of quantum mechanics. It is this behaviour that both motivates and enables our work using atoms to create, manipulate, and preserve quantum information encoded in light. Understanding the behaviour of light and atoms and the interactions between them will be necessary to interpreting the main work undertaken for this thesis, and so this first chapter is devoted to the quantum mechanical description of light and the interaction of light with atoms. For those less familiar with the mathematical framework, a more complete outline can be found in Lambropoulos and Petrosyan [1].

2.1 Quantum theory

Quantum mechanics is based on the description of physical objects as being in quantum states and how these states evolve with time. This is most familiar in assigning a wave function to the properties of an object. For example, an object represented by $|\psi\rangle$ existing in one continuous spatial dimension could be described by

$$|\psi\rangle = \int \psi(x)\,|x\rangle\,dx, \tag{2.1}$$

with the probability of measuring the object in a small region x' of

$$P(x') = \int_{x'} |\psi(x)|^2 dx. \tag{2.2}$$

The expectation value of the position operator $\hat{X}$, which performs $\hat{X}\,|\psi\rangle = x\,|\psi\rangle$ is then

$$\langle\psi|\,\hat{X}\,|\psi\rangle = \int_{-\infty}^{\infty} x|\psi(x)|^2 dx, \tag{2.3}$$

which can also be written $\left\langle \hat{X} \right\rangle$.

In the Schrödinger picture states evolve according to the equation

$$i\hbar \frac{d\left|\psi(t)\right\rangle}{dt} = \hat{H}\left|\psi(t)\right\rangle \tag{2.4}$$

where $\hat{H}$ is the Hamiltonian, the operator giving the total energy of the state. Energy is considered to be an observable. This places restrictions on the Hamiltonian as the expectation value is required to be real. The Hamiltonian written in its eigenbasis is

$$\hat{H} = \sum_{n}^{N} E_n \left|E_n\right\rangle\!\left\langle E_n\right| . \tag{2.5}$$

where E_n are the eigenvalues representing the allowed energies of the system and the terms $|E_n\rangle\langle E_n|$ From Eqs. (2.4) and (2.5), a state that is an eigenstate of the Hamiltonian will have a definite energy, and will not evolve apart from accumulating a phase.

2.1.1 Superposition and entanglement

Superposition and entanglement are two defining characteristics of quantum mechanics. These can be represented in the following manner.

The Hilbert space of the system defines the dimensions over which the states of the system exist. These dimensions are also called degrees of freedom. In a space represented by discrete basis states, a pure state can be represented as

$$\left|\psi\right\rangle = \sum_{j} c_j \left|j\right\rangle \tag{2.6}$$

where c_j are complex numbers, $|c_j|^2$ is the probability of finding the system in the state $|j\rangle$, and $\sum_j |c_j|^2 = 1$.

Superposition describes the ability for a system to to be in more than one basis state at the same time. This can be distinguished from the system being in a random choice of these states by the way the superposition affects measurements. For example, a single photon may take both paths through a double slit experiment; the two paths may generate an interference pattern at some subsequent position.

Regarding entanglement, consider a system consisting of an expanded Hilbert space $H_A \otimes H_B$, where $\otimes$ is the tensor product. A and B might describe different degrees of freedom of a system or different subsystems. A state in this expanded space can be represented as

$$\left|\Psi\right\rangle_{AB} = \sum_{i,j} c_{ij} \left|i\right\rangle_A \otimes \left|j\right\rangle_B . \tag{2.7}$$

The state is separable over the subspaces if the weights can be represented $c_{ij} = c_i^A c_j^B$. Otherwise the subsystems are said to be entangled.

2.1.2 Density matrix representation

In some interpretations of quantum mechanics, the universe is considered to be in a pure state. For the sake of understanding the behaviour of smaller parts of the universe, sys-

tems in classical distributions should be considered.

A system in a classical distribution of basis states $|\psi_n\rangle$ with probabilities for each state P_n can be written as the density operator,

$$\hat{\rho} = \sum_{n=1}^{N} P_n \left|\psi_n\right\rangle\!\left\langle\psi_n\right| . \tag{2.8}$$

The evolution of the density operator can be written

$$i\hbar\partial_t\hat{\rho} = [H, \hat{\rho}]. \tag{2.9}$$

Density matrix formalism is useful for understanding the results of measurements on subsystems. For example, measurements performed on the subsystem A of a system in the Hilbert space $H_A \otimes H_B$ will give results according to the reduced density matrix $\hat{\rho}_A$. The reduced density matrix is found by tracing over the rest of the system,

$$\hat{\rho}_A = \sum_j \langle j|_B \, \hat{\rho}_{AB} \, |j\rangle_B = \mathrm{Tr}_B[\hat{\rho}_{AB}]. \tag{2.10}$$

The trace also gives a neat representation for the expectation value of an observable,

$$\left\langle \hat{\mathcal{O}} \right\rangle = \mathrm{Tr}[\hat{\rho}\hat{\mathcal{O}}]. \tag{2.11}$$

The whole system may be in a pure state, but if the subsystem is not separable from the rest of the system then under local measurements the subsystem will be a mixed state. Apart from the interpretive questions this formalism raises, it also explains the requirement to ensure the measured degrees of freedom will not couple to their environments. Quantum information encoded as superposition and entanglement within a system will be lost if the system interacts with the environment and generates entanglement between the measured degrees of freedom and external degrees of freedom. The tendency of subsystems to be in mixed states also leads to the requirement for preparation steps to place an experimental system in an initially pure state.

2.1.3 Qubits

A qubit is essentially a two-state system. There will however be other degrees of freedom that allow the qubits in an expanded system to be distinguished from each other. In the case of photon-based qubits these extra degrees of freedom allow the qubit to be stored, retrieved, and moved around.

The qubit must interact with other qubits in most quantum computing schemes. In theoretical terms, this means constructing a Hamiltonian that depends on the states of two qubits. It is important that the additional degrees of freedom are not entangled during the interaction. In Part III we discuss how this can cause problems for photon-photon interactions, as the entanglement between additional degrees of freedom puts the combined qubit system into a mixed state instead of the desired pure state.

2.2 Quantum states of light

The following quantum description of light is sufficient to describe the interaction of light with atoms. The experimental sections in this thesis involve relatively simple states of light, but we can still test that the atomic memory does preserve the quantum characteristics of the input states. With a proper theoretical basis, it is also possible to predict how well the quantum information processes studied in this thesis perform.

2.2.1 The quantised electromagnetic field

The electric field of light propagating as a plane wave in free space along the dimension z can be described by the operator

$$\hat{\boldsymbol{E}}(z) = \boldsymbol{\epsilon}\sqrt{\frac{\hbar\omega}{4\pi c\epsilon_0 A}}\int_{-\infty}^{\infty} \mathrm{d}\omega \left[\hat{a}_\omega e^{i\omega z/c} + \hat{a}_\omega^\dagger e^{-i\omega z/c}\right], \tag{2.12}$$

where c is the vacuum speed of light, ϵ_0 is the vacuum permittivity, A is the transverse cross-section of the mode and $\boldsymbol{\epsilon}$ is the polarisation vector of the field. This representation describes an optical field existing as a continuum of angular frequencies ω. For the quantisation of the electromagnetic field, see [2].

The creation and annihilation operators $\hat{a}_w^\dagger$ and $\hat{a}_w$ obey the commutation relations,

$$\left[\hat{a}_\omega, \hat{a}_{\omega'}^\dagger\right] = \delta(\omega - \omega') \tag{2.13}$$

$$[\hat{a}_\omega, \hat{a}_{\omega'}] = 0 \tag{2.14}$$

$$\left[\hat{a}_\omega^\dagger, \hat{a}_{\omega'}^\dagger\right] = 0. \tag{2.15}$$

The Hamiltonian for the optical field is

$$\hat{H}_L = \int \hbar\omega \left(\hat{a}_\omega^\dagger \hat{a}_\omega + \frac{1}{2}\right) d\omega. \tag{2.16}$$

In all the cases we will consider, the optical fields used have a bandwidth much narrower than the frequency of the field. The field can then be described in terms of a carrier envelope in time and space,

$$\hat{\mathbf{E}}_p(z,t) = \boldsymbol{\epsilon}\sqrt{\frac{\hbar\omega_p}{4\pi c\epsilon_0 A}}\int_{\omega_p} \mathrm{d}\omega \left(\hat{a}_\omega e^{i\omega(t+z/c)} + \hat{a}_\omega^\dagger e^{-i\omega(t+z/c)}\right). \tag{2.17}$$

where ω_p is the central frequency of the field.

2.2.2 Optical quadratures

The electromagnetic field quadratures are represented in terms of the creation and annihilation operators:

$$\hat{x} = \frac{\hat{a}^\dagger + \hat{a}}{\sqrt{2}} \tag{2.18}$$

$$\hat{p} = \frac{i(\hat{a}^\dagger - \hat{a})}{\sqrt{2}}. \tag{2.19}$$

2.2.3 Fock states

Fock states, or number states, are states that have a well defined occupation number, such as number of photons in a mode. A number state is $|n\rangle$ where n is the occupation number. The state can be represented in terms of the creation operator acting on the vacuum, or zero occupation state as

$$|n\rangle = \frac{(\hat{a}^\dagger)^n}{\sqrt{n!}} |0\rangle . \tag{2.20}$$

The raising and lowering operators acting on Fock states give

$$\hat{a} |n\rangle = \sqrt{n} |n-1\rangle , \hat{a}^\dagger |n-1\rangle = \sqrt{n} |n\rangle . \tag{2.21}$$

The number operator acting on a Fock state gives

$$\langle n| \hat{a}^\dagger \hat{a} |n\rangle = n \langle n|n\rangle = n. \tag{2.22}$$

Number states do not have a well defined phase. However, two Fock states can have a well defined relative phase.

2.2.4 Coherent states

An ideal laser emits a coherent state of light. A coherent state can be represented as a superposition of Fock states,

$$|\alpha\rangle = e^{-\frac{|\alpha|^2}{2}} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle = e^{-\frac{|\alpha|^2}{2}} e^{\alpha \hat{a}^\dagger} |0\rangle . \tag{2.23}$$

The expectation value of the number of photons is $\langle n\rangle = \langle \hat{a}^\dagger \hat{a}\rangle = |\alpha|^2$. The variance in the number of photons detected is also equal to the number of photons. This gives Poissonian statistics in which the arrival time of photons at a detector is uncorrelated between the photons.

2.3 Measuring and characterising quantum states of light

The electric field of light can be detected down to the single photon level in a destructive measurement (such as with a photodiode). Each photon is absorbed in the detector, releasing an electron with probability equal to the 'quantum efficiency' of the detector. The electron contributes to a measurable photocurrent which has a variance given by the photon variance in the detection bandwidth, with some amount of added noise due to the detector electronics. This detection cannot directly detect the phase of the field.

2.3.1 Measurement techniques

Heterodyne measurement can be used to determine optical phase by combining the light at a beamsplitter with another optical field, as shown in Fig. 2.1. This second field is called a local oscillator (LO). The interference between the two fields at the beamsplitter outputs is used to determine both the phase and amplitude, with additional vacuum noise added due to measuring non-commutative operators. This technique can also be

used to reduce the effect of detector noise on the measurement, as the local oscillator field multiplies the signal field after the interference.

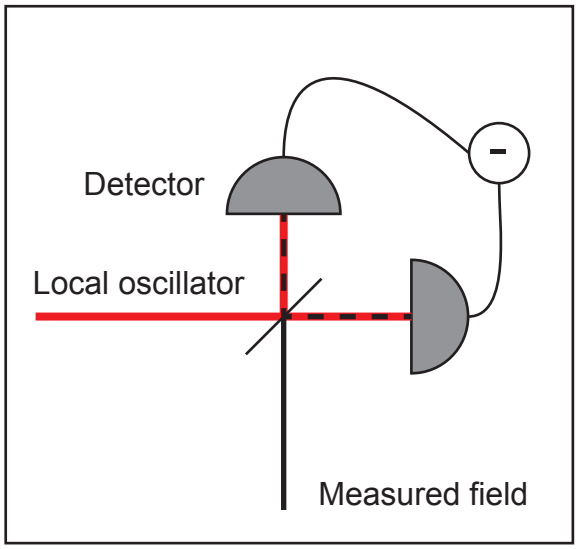

Fig. 2.1. In heterodyne detection, the optical field to be measured is combined at a beamsplitter with a bright classical field detuned from the measured field called a local oscillator. The resulting outputs are measured and the signals are subtracted. A beat note is produced in the resulting signal at twice the detuning, and dependent on the amplitude of both the local oscillator and measured fields. The phase of the measured field is then known in relation to the phase of the local oscillator.

2.3.2 Descriptions of quantum states of light

We have looked so far at pure states of light, but we need to be able to describe other states of light that we might come across. One form for describing mixed states of light is the Husimi Q function [3].

The Q function is a description of light in the basis of coherent states. It is based on the diagonal elements of the density matrix and is written as

$$Q(\alpha, \alpha^*) = \frac{1}{\pi} \langle \alpha | \rho | \alpha \rangle \tag{2.24}$$

For example, the Q function for a coherent state $\rho_0 = |\alpha_0\rangle\langle\alpha_0|$ is [4]

$$Q(\alpha, \alpha^*) = \frac{1}{\pi} |\langle \alpha_0 | \alpha \rangle|^2 = \frac{1}{\pi} \exp\left(-|\alpha - \alpha_0|^2\right). \tag{2.25}$$

Realistic states will have a broader distribution in the Q function, and so the distribution can be used to measure the noise added by a process. Other distributions such as the Wigner function can be arrived at from the Q function.

Heterodyne measurement with a sufficiently strong local oscillator is a direct sampling of the Q function [5]. Tomographic measurement of a state by repeated preparation and heterodyne measurement at all angles is therefore a useful measurement technique.

2.3.3 Characterising quantum processes

After measuring and characterising a state, we would like to know how effective the process generating the state is. For example, the output of the ideal quantum memory will reproduce the input state exactly at some later time. The effectiveness of the memory can be judged by how close the output state is to the input state.

Fidelity

A common measure of distance between states is fidelity. The fidelity of one state $|\Psi_O\rangle$ compared to the desired state $|\Psi_D\rangle$ is equal to the square of the overlap,

$$\mathcal{F} = |\langle\Psi_D|\Psi_O\rangle|^2. \tag{2.26}$$

The measure of fidelity for a mixed output state ρ_O and the desired state ρ_D is

$$\mathcal{F} = \mathrm{Tr}\left[\sqrt{\sqrt{\rho_O}\rho_D\sqrt{\rho_O}}\right]. \tag{2.27}$$

Some processes have different fidelities depending on the inputs. Calculating a measure such as fidelity across all combinations of inputs is called process tomography. By appropriately combining the effectiveness of each process, the accuracy of an entire operation can be estimated.

2.4 Atom-light interaction

Light interacts with atoms primarily by its electric field, which exerts a force on the charged particles of the atoms. The dipole interaction involving the relative motion of the positive and negative charges of the atom is the primary interaction describing the emission and absorption of light by atoms. While this theoretical approach focusses on atoms, it applies to some extent to any ensemble of emitters that interact with optical fields via the dipole interaction.

We start with the simplest form of the interaction: A two-level atom interacting with an optical field. For the following theory we rely primarily on the derivation by Daniel Steck [6].

2.4.1 Interaction of a two-level atom with an quantum optical field

We introduce an atom that has a ground state $|g\rangle$ and an excited state $|e\rangle$. The following sets of operators are useful for describing the state of the atom and its interactions with optical fields: The Pauli matrices,

$$\hat{\sigma}_x = |e\rangle\langle g| + |g\rangle\langle e| \tag{2.28}$$
$$\hat{\sigma}_y = i(|g\rangle\langle e| - |e\rangle\langle g|) \tag{2.29}$$
$$\hat{\sigma}_z = |e\rangle\langle e| - |g\rangle\langle g|\,, \tag{2.30}$$

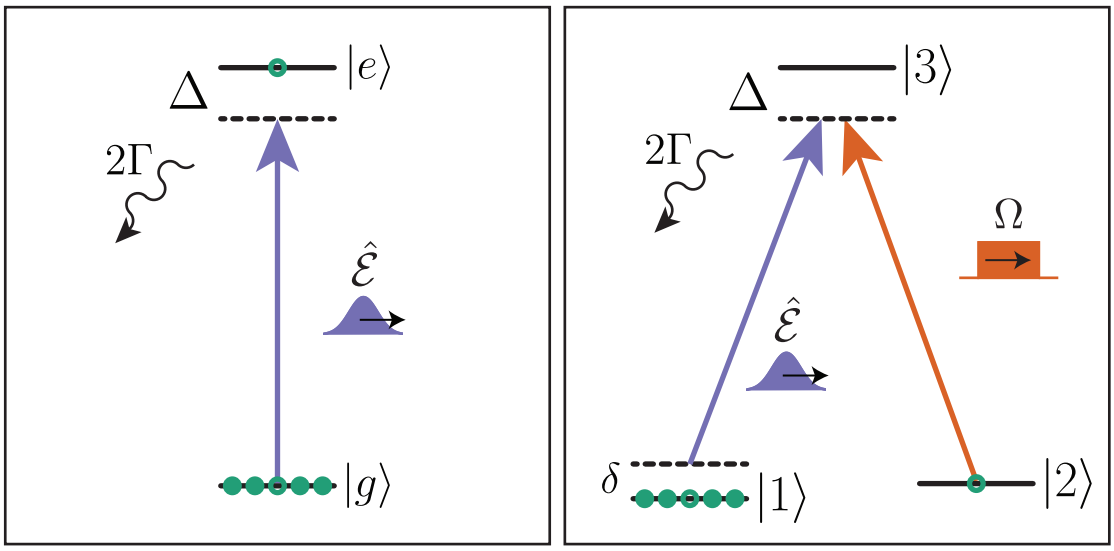

Fig. 2.2. Level schemes and detunings for two- and three-level atoms. On the left is an ensemble of two-level atoms interacting with a quantum field $\hat{\mathcal{E}}$. On the right is an ensemble of three-level atoms interacting with a quantum field $\hat{\mathcal{E}}$ and a classical field with Rabi frequency Ω.

and the atomic raising and lowering operators,

$$\hat{\sigma}_+ = \frac{1}{2}(\hat{\sigma}_x + i\hat{\sigma}_y) = |e\rangle\langle g| \tag{2.31}$$

$$\hat{\sigma}_- = \frac{1}{2}(\hat{\sigma}_x - i\hat{\sigma}_y) = |g\rangle\langle e|\,. \tag{2.32}$$

To find the interaction of a two-level atom with a light field with n photons, we can write the total Hamiltonian as the sum of the Hamiltonians of the atom, the light, and the coupling between the two,

$$H = H_A + H_L + H_{AL}. \tag{2.33}$$

The atomic Hamiltonian is simply the energy of the atomic states,

$$H_A = \frac{\hbar\omega_{eg}}{2}\hat{\sigma}_z, \tag{2.34}$$

where ω_{eg} is the resonant frequency for the transition. The energy is normalised so the zero energy point is halfway between the two atomic levels.

By treating the atom as a dipole, we find the energy for the interaction of the electric field of the light with the atom as

$$H_{AL} = -\hat{\boldsymbol{d}}.\hat{\boldsymbol{E}} \tag{2.35}$$

The dipole operator $\hat{\boldsymbol{d}}$ describes the orientation and size of the dipole formed by the atomic charges. The dot product in Eq. (2.35) projects the dipole onto the orientation of the electric field, giving the equivalent of the energy of a classical dipole in a field. The

rotational symmetry of the atom gives the average dipole as zero. That is, the quantity $\langle g|\hat{\boldsymbol{d}}|g\rangle = 0$. Dipole coupling is non-zero only between states of opposite parity, as the difference in phase of the two states breaks the symmetry of the action of the dipole operator. We assume this is the case for the atomic levels $|g\rangle$ and $|e\rangle$. We can then consider the dipole operator for the atom as only including the cross-terms:

$$\hat{\boldsymbol{d}} = \langle g|\hat{\boldsymbol{d}}|e\rangle\,|g\rangle\langle e| + \langle e|\hat{\boldsymbol{d}}|g\rangle\,|e\rangle\langle g| \tag{2.36}$$

We can also set $\langle g|\boldsymbol{d}|e\rangle\,|g\rangle\langle e| = \langle e|\hat{\boldsymbol{d}}|g\rangle\,|e\rangle\langle g|$ as these are only distinguished by a phase term. We call this quantity the coupling strength g, giving

$$\hat{\boldsymbol{d}} = g(\hat{\sigma}_+ + \hat{\sigma}_-) \tag{2.37}$$

This gives the atom-light part of the Hamiltonian

$$H_{AL} = -\hat{\mathbf{d}}.\hat{\boldsymbol{E}} \tag{2.38}$$

$$= \hbar g(\hat{\sigma}_+ + \hat{\sigma}_-)(\hat{a}^\dagger + \hat{a}). \tag{2.39}$$

The interaction of a quantum state of light with an atom by the dipole force is then described by the complete Hamiltonian, known as the Jaynes-Cummings Hamiltonian,

$$\hat{H} = \frac{\hbar\omega_{eg}}{2}\hat{\sigma}_z + \hbar\omega\hat{a}^\dagger\hat{a} + \hbar g(\hat{\sigma}_-\hat{a}^\dagger + \hat{\sigma}_+\hat{a}). \tag{2.40}$$

The term $\hat{\sigma}_+\hat{a}$ corresponds to the annihilation of a photon and the raising of the atom from the ground to the excited state, while $\hat{\sigma}_-\hat{a}^\dagger$ corresponds to the creation of a photon and the lowering of the atom from the excited to the ground state. The terms $\hat{\sigma}_+\hat{a}^\dagger$ and $\hat{\sigma}_-\hat{a}$ do not conserve energy, averaging to zero over the time scales of experimental relevance, and are therefore neglected. The more rigorous method for removing these terms is known as the rotating wave approximation.

2.4.2 Dressed states

Where there is strong coupling between atom and light, the system can no longer be considered as separate states of the two. By expressing the optical field as a number state $|n\rangle$, the Jaynes-Cummings Hamiltonian gives

$$\langle g, n+1|H|e, n\rangle = \hbar g\sqrt{n+1}. \tag{2.41}$$

with g the coupling strength. Energies of the uncoupled systems are

$$E_{e,n} = \langle e, n|\hat{H}|e, n\rangle = \hbar(\omega_{eg}/2 + n\omega) \tag{2.42}$$

$$E_{g,n+1} = \langle g, n+1|\hat{H}|g, n+1\rangle = \hbar(-\omega_{eg}/2 + (n+1)\omega). \tag{2.43}$$

Introducing the detuning $\Delta = \omega_{eg} - \omega$ and diagonalising the Hamiltonian gives eigenvalues

$$E_n^\pm = \hbar\Delta/2 \pm \hbar\Omega_n, \tag{2.44}$$

with corresponding eigenstates

$$|\pm_n\rangle = \frac{1}{\sqrt{N_\pm}}(\Omega_n \mp \Delta/2)\,|g, n+1\rangle \pm g\sqrt{n+1}\,|e, n\rangle\,, \tag{2.45}$$

where the Rabi frequency $\Omega_n = \sqrt{g^2(n+1) + (\Delta/2)^2}$ and $N_\pm = g^2(n+1) + (\Omega_n \mp \Delta/2)^2)$. These eigenstates are called the dressed states.

2.4.3 Interaction of two-level atoms with a classical field

We can also consider the case of an atom interacting with a classical field, where the field is described by

$$\boldsymbol{E}(z,t) = \boldsymbol{\epsilon}\mathcal{E}(t)\cos\left[\omega(t - z/c)\right] = \frac{\boldsymbol{\epsilon}}{2}\mathcal{E}(t)\left(e^{-i\omega t} - e^{i\omega t}\right). \tag{2.46}$$

We can define the Rabi frequency in this case as

$$\Omega(t) = \frac{1}{2\hbar}(\hat{\boldsymbol{d}}.\boldsymbol{\epsilon})\mathcal{E}(t) \tag{2.47}$$

and similarly to the case of the quantum field, the complete Hamiltonian is (ignoring the light Hamiltonian)

$$\hat{H} = \hbar\Delta\hat{\sigma}_z - \frac{i\hbar\Omega}{2}(\hat{\sigma}_-\hat{a}^\dagger + \hat{\sigma}_+\hat{a}). \tag{2.48}$$

The evolution of an atom can then be written in terms of the atomic operators as

$$\begin{aligned}
\partial_t \langle\hat{\sigma}_x\rangle &= \Delta\,\langle\hat{\sigma}_y\rangle \\
\partial_t \langle\hat{\sigma}_y\rangle &= -\Delta\,\langle\hat{\sigma}_x\rangle - \Omega\,\langle\hat{\sigma}_z\rangle \\
\partial_t \langle\hat{\sigma}_z\rangle &= \Omega\,\langle\hat{\sigma}_y\rangle\,.
\end{aligned} \tag{2.49}$$

2.4.4 Bloch sphere representation

The expectation values of the Pauli operators are a useful basis for visualising the evolution of the atom. By assigning to each expectation value $\left\langle\hat{\sigma}_{(x,y,z)}\right\rangle$ a corresponding spatial dimension (x, y, z) the state of the atom can be projected onto a sphere. This is the Bloch sphere representation of the atom, shown in Fig. 2.3.

According to Eqs. (2.49), the application of a classical field on resonance for a duration $\Omega t = \pi$ drives the atom from $\langle\hat{\sigma}_z\rangle = -1$ into $\langle\hat{\sigma}_y\rangle = 1$ and through to $\langle\hat{\sigma}_z\rangle = 1$. This pulse producing a rotation of π about the x-axis is a π pulse.

2.4.5 Spontaneous decay of atoms

Atoms couple to the environment by the spontaneous emission of radiation. This interaction can be accounted for by assuming the environment is a reservoir of vacuum modes. The environment is treated as Markovian; it contains no information about previous atomic decays. The coupling of the atom to the reservoir can be modelled with the master equation formalism [1]. This formalism reproduces the density matrix evolution based on the averaging over ensembles of the spontaneous emission process, or

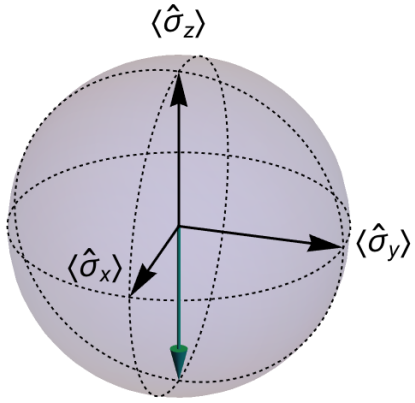

Fig. 2.3. Bloch sphere representation of a two-level atom with $\langle\hat{\sigma}_z\rangle = -1$. The green arrow represents the atomic state.

equivalently the unknown state of a single atom based on the same coupling. The master equation describes the evolution of the density matrix of the atom as

$$\partial_t\hat{\rho} = -\frac{i}{\hbar}[\hat{H},\hat{\rho}] + 2\Gamma\mathcal{D}[\sigma]\hat{\rho} \tag{2.50}$$

with the Lindblad superoperator

$$\mathcal{D}[\sigma]\hat{\rho} = \sigma_-\hat{\rho}\sigma_+ - \frac{1}{2}(\sigma_+\sigma_-\hat{\rho} + \hat{\rho}\sigma_+\sigma_-). \tag{2.51}$$

The evolutions of the density operator elements are

$$\partial_t\hat{\rho}_{gg} = 2\Gamma\hat{\rho}_{ee} - i(\hat{\rho}_{eg}\Omega^* + \hat{\rho}_{ge}\Omega) \tag{2.52}$$
$$\partial_t\hat{\rho}_{ee} = -2\Gamma\hat{\rho}_{ee} + i(\hat{\rho}_{eg}\Omega^* + \hat{\rho}_{ge}\Omega) \tag{2.53}$$
$$\partial_t\hat{\rho}_{eg} = -(\Gamma - i\Delta)\hat{\rho}_{eg} + i\Omega(\hat{\rho}_{gg} - \hat{\rho}_{ee}). \tag{2.54}$$

For very weak driving fields, the spontaneous decay ensures the atom is predominantly in the ground state and we can assume $\hat{\rho}_{gg} = 1$ and $\hat{\rho}_{ee} = 0$. This allows the atomic state to be described by

$$\partial_t\hat{\rho}_{eg} = -(\Gamma + i\Delta)\hat{\rho}_{eg} + i\Omega. \tag{2.55}$$

2.4.6 Collective atomic operators for two-level ensembles

To extend the theory to an ensemble of atoms experiencing the same field, we can define collective operators for the atoms at a particular spatial position z.

$$\hat{\sigma}_{\mu\mu}(z,t) = \frac{1}{N_z}\sum_{n=1}^{N_z}\hat{\sigma}^n_{\mu\mu}(t) \tag{2.56}$$

$$\hat{\sigma}_{eg}(z,t) = \frac{1}{N_z}\sum_{n=1}^{N_z}\hat{\sigma}^n_{eg}(t)e^{-i\omega_{eg}(t-z_n/c)} \tag{2.57}$$

Similarly, a slowly varying operator can be defined for the optical field:

$$\hat{\mathcal{E}}(z) = \sqrt{\frac{L}{2\pi c}}e^{i\omega_{eg}(t-z.c)}\int_{\omega_{eg}} d\omega \hat{a}_\omega(t)e^{i\omega z/c}. \tag{2.58}$$

The equation for light travelling through the ensemble can be found by

$$(\partial_t + c\partial_z)\hat{\mathcal{E}}(t,z) = -\frac{i}{\hbar}[\hat{H},\hat{\mathcal{E}}(t,z)] \tag{2.59}$$

$$= igN_z\hat{\sigma}_{eg} \tag{2.60}$$

Free induction decay

In an ensemble of atoms, the resonant frequency can be different for each atom due to differences in the local environment. This is known as inhomogeneous broadening. We can see from Eqs. (2.49) that a collection of atoms with different detunings Δ initially placed in $\langle\hat{\sigma}_y\rangle = 1$ will, ignoring decay to the ground state, eventually end up in different positions around the Bloch sphere equator. The length of the vector for the collective operator is the sum of atomic coherences with the range of different phases, and so will shorten as the ensemble dephases. The time this takes to occur can be equated with the T_2 lifetime:

$$\hat{\sigma}_{eg}(t) = \hat{\sigma}_{eg}(t_0)e^{-(t-t_0)/T_2} \tag{2.61}$$

Spin echo

A pulse on resonance can excite an ensemble of atoms into a superposition of ground and excited states. However, inhomogeneous broadening will cause the phase of each atomic superposition to precess at different rates. The states will smear out around the Bloch sphere as the phase of the superposition becomes different for each spin. A π pulse applied after time τ will rotate each state about the axis $\hat{y}$, effectively flipping the order of the states in their precession about the $\hat{z}$ axis. The states will rephase as the higher frequency atoms have been placed behind the lower frequency atoms in phase. The original pulse will be re-emitted at time 2τ. This process is shown on the Bloch sphere in Fig. 2.4.

The spin echo depends on the inhomogeneous broadening being constant over the time between the input pulse and the recalled pulse. The spin echo technique can be extended to a technique known as dynamical decoupling, where the π pulses are applied

more quickly than the broadening can change due to varying external fields. This allows compensating for these sources of broadening without having to measure them.

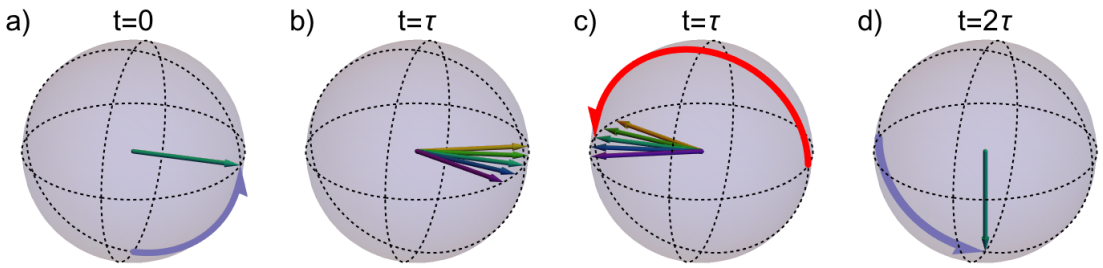

Fig. 2.4. Bloch sphere representation of a spin echo process. a) At time $t = 0$ the ensemble is moved to $\langle\hat{\sigma_y}\rangle = 1$ by a $\pi/2$ pulse (represented by a blue arrow). b) After some time the inhomogeneous broadening spreads the ensemble out along the equator of the Bloch sphere (the different coloured arrows represent different detunings). c) At time $t = \tau$ a π pulse (represented by a red arrow) flips the ensemble about the x-axis. This flips the relative positions of the emitters, and in d) at time $t = 2\tau$ the emitters have moved back into phase, re-emit a $\pi/2$ pulse, placing them back in the ground state $\langle\hat{\sigma_z}\rangle = -1$.

2.4.7 Interaction of three-level atoms with optical fields

For the three-level theory we also rely on the derivation in [7]. Use of three-level atoms typically involves driving the atoms with a bright classical field to modify the interaction between a weak field and the atoms. The operators for a weak quantum field with a co-propagating classical control field are then

$$\hat{\mathbf{E}}_p(z) = \epsilon_p \left(\frac{\hbar\omega_{p+}}{4\pi c\epsilon_0 A}\right)^{1/2} \int_{\omega_p} \mathrm{d}\omega \left(\hat{a}_\omega e^{i\omega z/c} + \hat{a}_\omega^\dagger e^{-i\omega z/c}\right)$$
$$\mathbf{E}_c(z) = \epsilon_c \mathcal{E}_c(t - z/c)\mathrm{Cos}[\omega_c(t - z/c)].$$

where we assume optical modes that each exist in a small bandwidth around a carrier frequency given by $\omega_p = \omega_{13} + \Delta + \delta$ and $\omega_c = \omega_{23} + \Delta$. The frequency ω_{13} is the resonant frequency for the transition $|1\rangle \rightarrow |3\rangle$, and similarly for ω_{23}. We use the same correspondence between classical field and Rabi frequency as in Eq. (2.47). We assume the probe field only couples with the transition $|1\rangle \rightarrow |3\rangle$ and the control field with the transition $|2\rangle \rightarrow |3\rangle$. The interaction part of the Hamiltonian is then

$$\hat{V} = -\hbar \sum_{i=1}^{N} \left[\Omega_c(t - z_i/c)e^{-i\omega_{c+}(t-z_i/c)}\hat{\sigma}_{23}^i + g\left(\frac{L}{2\pi c}\right)^{1/2} \int_{\omega_{p+}} \mathrm{d}\omega \hat{a}_\omega e^{i\omega z/c}\hat{\sigma}_{13}^i + \mathrm{H.C.}\right] \tag{2.62}$$

Collective atomic operators for three-level atoms

We can again assume that the atoms at any particular location all behave similarly, and so can be treated as smaller ensembles. We can remove optical frequencies from the

equations by defining the following slowly varying collective operators:

$$\begin{aligned}
\hat{\sigma}_{\mu\mu}(z,t) &= \frac{1}{N_z}\sum_n^{N_z}\hat{\sigma}_{\mu\mu}^n(t) \\
\hat{\sigma}_{32}(z,t) &= \frac{1}{N_z}\sum_n^{N_z}\hat{\sigma}_{32}^n(t)e^{-i\omega_c(t-z_n/c)} \\
\hat{\sigma}_{31}(z,t) &= \frac{1}{N_z}\sum_n^{N_z}\hat{\sigma}_{31}^n(t)e^{-i\omega_p(t-z_n/c)} \\
\hat{\sigma}_{21}(z,t) &= \frac{1}{N_z}\sum_n^{N_z}\hat{\sigma}_{21}^n(t)e^{-i(\omega_p-\omega_c)(t-z_n/c)} \\
\hat{\mathcal{E}}(z,t) &= \sqrt{\frac{L}{2\pi c}}e^{i\omega_p(t-z/c)}\int_{\omega_p}\mathrm{d}\omega\hat{a}_\omega(t)e^{i\omega z/c}
\end{aligned} \tag{2.63}$$

by assuming that each slice dz of the ensemble contains a number of atoms $N_z \gg 1$. The commutators for the collective operators are

$$\begin{aligned}
[\hat{\sigma}_{\mu\nu}(t),\hat{\sigma}_{\alpha\beta}(t)] &= \delta_{\nu\alpha}\hat{\sigma}_{\mu\beta}(t) - \delta_{\mu\beta}\hat{\sigma}_{\alpha\nu}(t) \\
\left[\hat{\mathcal{E}}_\pm(t),\hat{\mathcal{E}}_\pm^\dagger(t)\right] &= 1.
\end{aligned}$$

The slowly varying operators are inserted into $\hat{V}$, giving a Hamiltonian

$$\begin{aligned}
\hat{H} = \int d\omega\hbar\omega\hat{a}_\omega^\dagger\hat{a}_\omega - \hbar\omega_p\frac{1}{L}\int_0^L dz\hat{\mathcal{E}}^\dagger\hat{\mathcal{E}}+ \\
\int_0^L \mathrm{d}z\hbar N_z\times\left(\Delta\hat{\sigma}_{33} - \left[\Omega_c(t-z/c)\hat{\sigma}_{32}+\right.\right. \\
\left.\left.\Omega_{c+}(t-z/c)\hat{\sigma}_{32}e^{-i[\omega_c(t-z/c)]} + g\hat{\mathcal{E}}\hat{\sigma}_{31} + \mathrm{H.c.}\right]\right),
\end{aligned} \tag{2.64}$$

where the (z,t) dependence of the operators is omitted for readability.

Optical Bloch equations

The evolution of the operators can again be found by taking the master equation approach,

$$i\hbar\partial_t\rho' = [H,\rho'] + \mathcal{L}\rho', \tag{2.65}$$

where $\mathcal{L}$ is the appropriate superoperator for three-level atoms.

The spontaneous decay term Γ again appears, behaving in the same manner as for the two-level atom calculation. The transition between ground states $|1\rangle \leftrightarrow |2\rangle$ is dipole forbidden, and the spontaneous emission from one state to the other is negligible. However interactions affecting the phase of the coherence still contribute to a decay of the ground state coherence $\hat{\sigma}_{12}$. The decay term γ can then also be extracted from the master equation. The magnitude of this decay depends on dephasing interactions and atomic

motion and can be measured as an experimental parameter.

In the case of a weak probe field, only the following three equations are needed to completely describe the system:

$$\partial_t \hat{\sigma}_{13} = -(\Gamma + i\Delta)\hat{\sigma}_{13} + ig\hat{\mathcal{E}} + i\Omega_c \hat{\sigma}_{12} \tag{2.66}$$

$$\partial_t \hat{\sigma}_{12} = -(\gamma + i\delta)\hat{\sigma}_{12} + i\Omega_c^* \hat{\sigma}_{13} \tag{2.67}$$

$$(\partial_t + c\partial_z)\hat{\mathcal{E}} = igN_z \hat{\sigma}_{13}. \tag{2.68}$$

2.4.8 Linear susceptibility

Assuming the field that couples to the populated level is weak (we can ignore the change in population of the atomic levels), we can solve for the steady state coherences. This then allows a straightforward description of how light interacts with the ensemble. This gives the linear susceptibility of the ensemble [8],

$$\chi^{(1)}(\delta) \propto \frac{4\delta(|\Omega_c|^2 - 4\delta\Delta) - 4\Delta\gamma^2}{||\Omega_c|^2 + (\Gamma + i2\Delta)(\gamma + i2\delta)|^2} + i\frac{8\delta^2\Gamma + 2\gamma(|\Omega_c|^2 + \gamma\Gamma}{||\Omega_c|^2 + (\Gamma + i2\Delta)(\gamma + i2\delta)|^2} \tag{2.69}$$

The notation is slightly different in the above equation: The control field detuning is set to Δ and the probe field detuning is set to δ.

2.5 Electromagnetically-induced transparency

A probe optical field resonant with a populated atomic transition is normally strongly absorbed by those atoms. However, a strong optical field resonant with a second (unpopulated) transition to the same excited state interferes with the absorption in a process known as electromagnetically-induced transparency (EIT), allowing the first field to pass through in a narrow bandwidth transparency window. The transparency appears due to a splitting of the transition caused by the control field. This transparency window causes strong dispersion and a reduced group velocity for the probe field. This relationship is shown in the susceptibility plot in Fig. 2.5. Light with a reduced group velocity due to the EIT effect is known as slow light.

A probe pulse travelling through an EIT medium is no longer entirely optical and becomes partly atomic excitation. The probe and control fields produce a coherent excitation of the medium into the auxiliary state. This is referred to as an atomic coherence, or spinwave (the term spinwave is derived from the fact that the auxiliary state and ground states are generally different spin states in the same orbital - while wave signifies a coherence.) The combination of light and atomic polarisation is a polariton. The propagation of the light in EIT can be described by the equation [9]

$$\left(\frac{\partial_t}{c} + \partial_z\right)\hat{E}(z,t) = -\frac{d}{\Omega(t)}\partial_t \frac{\hat{E}(z,t)}{\Omega(t)} \tag{2.70}$$

this gives a group velocity of the light of $v_g = \left(1/c + d/(L\Omega^2)\right)^{-1}$.

With sufficient optical depth, the sent probe pulse may fit entirely within the medium as a polariton. When the control field is adiabatically switched off, the probe field is entirely mapped onto the spinwave. There is no optical component, but the situation is

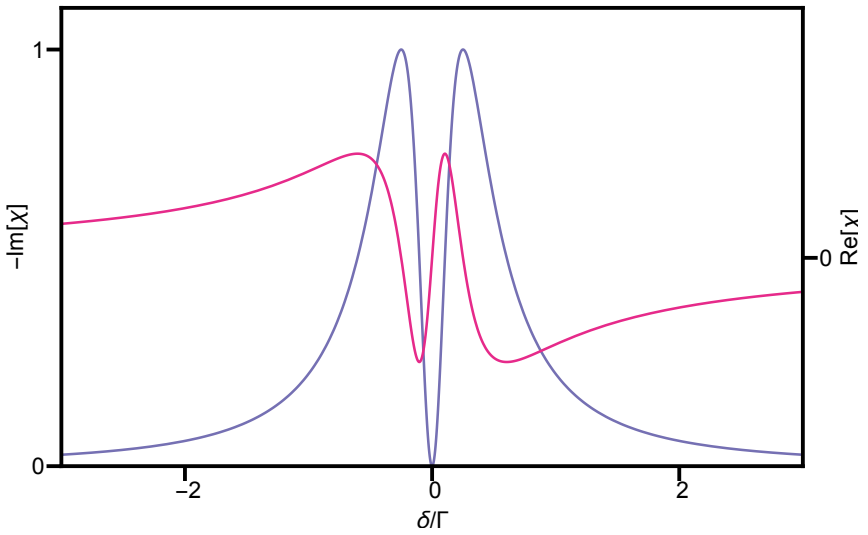

Fig. 2.5. Real (purple) and imaginary (blue) parts of susceptibility as the probe field frequency is scanned. Control field is resonant to $|2\rangle \rightarrow |3\rangle$ and control field Rabi frequency is equal to Γ. The ground state decoherence γ=0. The steep positive slope in $Re[\chi]$ corresponding with the minimum in $-Im[\chi]$ is the EIT region.

often referred to as 'stopped light'. The probe field can be retrieved from the medium by switching the control field back on.

2.6 Raman interaction

For a large one-photon detuning $\Delta \gg \Gamma$, the interaction on two-photon resonance becomes absorptive. The bandwidth of the absorptive feature can be much narrower than the one-photon linewidth Γ, allowing long lifetime memories that still take advantage of the strong excited state coupling. A narrow absorption due to a detuned control field is shown in Fig. 2.6.

We can also rewrite the equations in terms of parameters that are directly measurable from experiment. The collective operators are renormalised by defining $\hat{S} = \sqrt{N}\hat{\sigma}_{12}$ and $\hat{P} = \sqrt{N}\hat{\sigma}_{13}$. The optical field is renormalised to $\hat{\mathcal{E}} \rightarrow \sqrt{c/\Gamma}\hat{\mathcal{E}}$.

A change of frame is made, $z' = z + ct$, accounting for the finite speed of light and giving

$$\partial_{z'}\hat{\mathcal{E}} = i\frac{\sqrt{d}}{L}\hat{P} \tag{2.71}$$

The spatial coordinate z' is then replaced with a normalised position ξ which runs from 0 to 1. The equations can then be written in terms of the optical depth $d = g^2N/(\Gamma c)$, giving

$$\begin{aligned}
\partial_t\hat{P} &= -(\Gamma + i\Delta)\hat{P} + i\sqrt{d}\Gamma\hat{\mathcal{E}} + i\Omega_c\hat{S} \\
\partial_t\hat{S} &= -(\gamma + i\delta)\hat{S} + i\Omega_c^*\hat{P} \\
\partial_\xi\hat{\mathcal{E}} &= i\sqrt{d}\hat{P}.
\end{aligned} \tag{2.72}$$

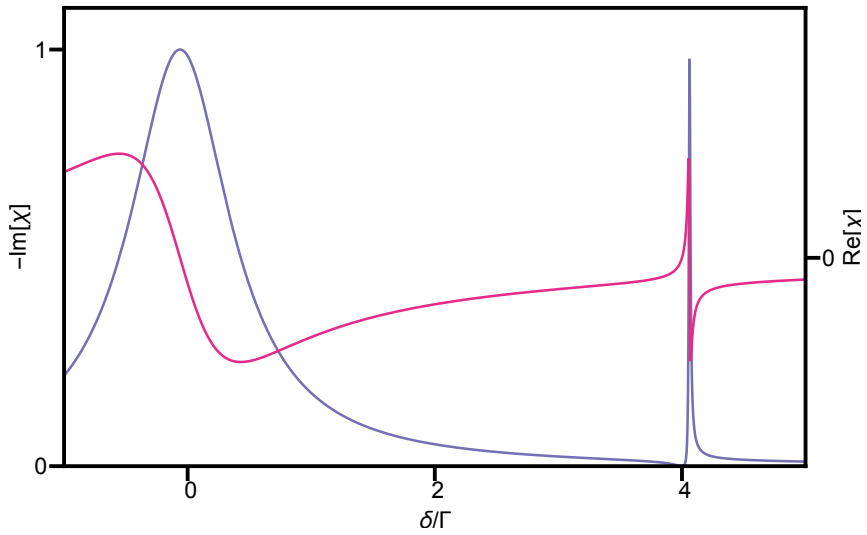

Fig. 2.6. Real (purple) and imaginary (blue) parts of susceptibility as the probe field frequency is scanned. Control field is detuned $\Delta = 4\Gamma$ to $|2\rangle \rightarrow |3\rangle$ and control field Rabi frequency is equal to 2Γ while the ground state decoherence γ=0. The absorption peak is slightly shifted by the AC-Stark effect. The asymmetry of the absorption peak reduces with larger $|\Delta|$.

The excited state may be adiabatically eliminated by assuming that its evolution is dominated by the detuning $\Delta \gg \partial_t \hat{P}$ so that

$$\hat{P} = i\left(\sqrt{d}\Gamma\hat{\mathcal{E}} + \Omega_c\hat{S}\right)/\left(\Gamma + i\Delta\right). \tag{2.73}$$

The equations of motion can now be expressed as

$$\partial_t\hat{S} = -(\gamma' + i\delta')\hat{S} + i\sqrt{d}\Gamma\left(\frac{\Omega_c^*}{\tilde{\Delta}}\hat{\mathcal{E}}\right) \tag{2.74}$$

$$\partial_\xi\hat{\mathcal{E}} = i\left(d\frac{\Gamma}{\tilde{\Delta}}\hat{\mathcal{E}} + \sqrt{d}\frac{\Omega_c}{\tilde{\Delta}}\hat{S}\right) \tag{2.75}$$

where

$$\delta' = \delta - \frac{|\Omega_c|^2\Delta}{\Gamma^2+\Delta^2},\ \gamma' = \gamma + \Gamma\frac{|\Omega_c|^2}{\Gamma^2+\Delta^2},\ \tilde{\Delta} = \frac{\Delta^2+\Gamma^2}{\Delta+i\Gamma}.$$

The optical depth for the relevant transition is measured by sending a weak probe field in the absence of a control field. The expected absorption is calculated by a steady state approximation for $\hat{P}$ and the assumption that no atoms are moved from the ground state. The expected absorption is then given by

$$\log\left(\frac{|E_{in}|}{|E_{out}|}\right) = \frac{d}{1+\Delta^2/\Gamma^2} \tag{2.76}$$

Dephasing and decay

The lifetime of a state on which information is stored fundamentally limits the storage time, as information is lost as the state decays. This lifetime is the T_1 lifetime. Between

ground states in atomic memories this lifetime is generally long enough to be ignored.

In a coherent memory, information is not just stored as the population of a state, but the relative phase between populations. The element ρ_{21} has a phase associated with it. Any uncontrolled inhomogeneous broadening of the populations causes a loss of this phase information. This restricts T_2 lifetime. T_2 lifetime also accounts for other physical mechanisms beyond broadening that cause the loss of this phase information, such as atomic motion.

Spinwave momentum

The spinwave has momentum as it is generated by absorbing a photon of one momentum and emitting one at another momentum. Higher momentum $k = 2\pi/\lambda_{SW}$ corresponds to shorter wavelength of the spinwave, λ_{SW}.

The spinwave momentum affects atom light interaction. Atomic motion can wash out the spinwave, as atoms carrying different phase information end up in the same place. Also, any change in the frequency or direction of the fields between generating the spinwave and recalling it will change the phase matching condition; The phase of the generated field at each point must match the phase difference between the control field and the spinwave. This new phase matching condition may affect the efficiency of the recall process.

AC-Stark shift

The term $\frac{|\Omega_c|^2\Delta}{\Gamma^2+\Delta^2}$ in Eq. (2.76) describes the shift in the energy separation of the atomic levels. This allows the use of bright optical fields to change the atomic energy levels, affecting the interaction of other light fields with the atoms or the evolution of states stored in atomic excitations. This effect is the basis of several schemes discussed in Part III to allow interactions between single photon states.

Part I

Atom-optic memories

Chapter 3

Literature review on optical memories and coherent optical processing

> Light thinks it travels faster than anything but it is wrong. No matter how fast light travels, it finds the darkness has always got there first, and is waiting for it.
>
> Terry Pratchett, *Reaper Man*

Light is a resource for quantum computing and communication. Memories that preserve the quantum information carried by the light are necessary for synchronisation and storing states until they are needed, multiplexing sources to create multi-photon states [10], and many more uses. Optical quantum memories have been pursued for the last two decades. Many useful platforms have been developed, with some on the verge of being put to use in real world quantum communication and computation applications.

The same platforms used to store quantum states can also be modified to allow for further processing of those states. In addition, the same linear properties that allow for high fidelity storage are also capable of enhancing nonlinear interactions between states of light. These interactions are necessary for building deterministic gates for optical quantum states, potentially allowing for quantum computation based on light. This possibility will be explored in the last part of this thesis.

In this chapter we review the literature on optical memories, focussing on ensemble-based memories. While most of the ensemble memories use atoms, some use ions in crystalline lattice and still others use colour centres in diamond. For this reason we will use the term 'emitter' to refer to a single element of an ensemble.

An optical quantum memory should deterministically store an optical quantum state and retrieve it on demand. For a state encoding a qubit, the memory must preserve the degree of freedom over which the qubit is encoded. The degree of freedom is normally chosen so that it does not affect the operation of the memory. We will generally assume this requirement is satisfied in favour of considering the overall preservation of a general optical quantum state.

The optical depth (OD) is an important resource for ensemble memories. The OD depends on the number of emitters and ultimately determines how much of the light is absorbed. In three-level systems, the OD also sets the branching ratio between light mapped into the memory and light incoherently scattered by the optical transition [11].

The superposition of emitters coherently emitting light allows the enhancement of the efficiency of coherent ensemble memories over single emitters.

The OD has further uses in multi-mode memories. The OD can be divided up so the different modes can all be stored efficiently. For temporally separated modes this allows a high time-bandwidth product. This product is another important property of memories, improving the usefulness of the memory for many of its applications.

3.1 Optical cavities and delay lines

The simplest quantum memory for light is a loop of optical fibre. A loss of about 0.15 dB/km is currently possible in optical fibre, limiting storage times to tens of microseconds. The delay time for a fibre loop is fixed and modifications are necessary to build an on demand memory. Some of these schemes have been demonstrated in free space memories.

This concept was demonstrated as a free space optical circulator [12] storing multiple time-binned photons to produce a reliable single photon source. Polarisation control was used to select when a photon was released from the device. In combination with a single-photon source, the device produced single photons more reliably than a similar non-multiplexed source, and could be significantly improved with state of the art components.

In another example, photon pairs were generated inside a cavity [13]. By dynamically controlling an adjacent shutter cavity, the heralding photon could first be released, and the other photon could be released up to 300 ns later with 50% single photon purity.

3.2 Ensemble memory platforms

For the rest of this chapter we will focus on ensemble memories. Before introducing the various memory schemes it is worth briefly explaining the main platforms in which the schemes are implemented.

Dilute atomic vapours make good ensembles, as each atom behaves basically identically. Achieving high optical depths can require high temperatures to increase the atomic vapour pressure. This raises the problem of Doppler broadening and atomic motion more generally. Atomic motion can be reduced by also including a noble gas such as xenon or krypton. Collisions with the noble gas reduce the mean free path of the ensemble atoms, reducing the diffusion of the atoms through the memory.

Cold and dense atomic vapours can be produced by use of a magneto-optic trap (MOT). Velocity-dependent scattering of light from atoms in combination with a magneto-optic potential is used to cool and trap atoms. The lifetime of these memories is generally limited by the cloud temperature as the trap must be turned off during storage. Some optical lattice schemes have been developed that allow the trapping to remain on during storage, greatly increasing the potential lifetime [14, 15].

In solid state systems the emitters are confined in a material lattice and motion is not a problem. However, solid state systems generally need preparation in order to produce a useful ensemble as lattice defects and impurities put strain on nearby lattice positions, generating large inhomogeneous linewidths on the transitions. Memory schemes such as atomic frequency comb take advantage of these naturally broadened transitions. Spin exchange between emitters and the lattice is another source of loss due to the more complex

electronic environment.

The inhomogeneous linewidth can be manipulated by 'hole burning'. This involves optically pumping emitters within a frequency bandwidth into other states and then moving narrower bandwidth groups of atoms back to create a homogeneously broadened ensemble. This is limited by the homogeneous linewidth of the atoms, the reliability of the optical pumping, and any relaxation of atoms back to the original state. Recent work in solid state has produced T_2 times of hours [16]. Quantum memories are also possible at telecom wavelengths with erbium doped crystals [17].

In solid state systems the emitters are not necessarily atoms or ions. An optical pulse was converted to phonons in diamond, allowing a THz bandwidth optical memory [18]. Nitrogen-vacancies in diamond have also been proposed as suitable ensembles [19].

3.3 Electromagnetically induced transparency

Electromagnetically induced transparency (EIT) was the first scheme to be proposed as an optical quantum memory, and the idea has been pursued since the time of the landmark EIT paper demonstrating light slowed to a group velocity of 17 m/s [20].

In 2001, the storage of an optical pulse was demonstrated by turning off the coupling field while a pulse was travelling through the memory [21]. The method of stopping light with EIT was analysed as a quantum memory [22]. Reviews of the earlier work in EIT can be found in [8, 23].

There have been many demonstrations of EIT as a quantum memory. For example, the correlations between DLCZ (see Sec. 3.6) read and write photons were preserved while also storing the write photon in a cold atomic ensemble [24]. This was also demonstrated in a warm rubidium vapour [25]. With an increased optical depth in cold atoms, 49% storage efficiency of single photons was achieved [26].

The current state of EIT as a quantum memory includes attempts to achieve high optical depth for a large delay-bandwidth product, dealing with problems such as absorption caused by the more complicated level schemes of real atoms, and eliminating ground-state dephasing and four-wave mixing to allow efficient and high fidelity memory.

A proposal for distinguishing between EIT and Autler-Townes splitting [27] was implemented in cold caesium atoms [28]. The test allows for precise measurement of some of the system properties important for implementing a memory.

In experiments focussing on high efficiency, 78% storage and recall efficiency is achieved, and a storage delay of 74 pulse widths at 50% efficiency [29]. A delay of a pulse width without storage was reported in 2014, along with a study of the different contributions of four wave mixing in the two species ^{87}Rb and ^{85}Rb [30].

Storage of a classical image was demonstrated in a solid state memory with a lifetime of one minute [31]. This long lifetime was achieved with the help of ensemble preparation and dynamical decoupling optimised by an evolutionary algorithm.

Beyond correlated photons generated with DLCZ, other quantum states of light have been stored with EIT. Several demonstrations exist of squeezed state storage [32, 33, 34, 35], where the squeezing was preserved upon retrieval of the light.

In 2008, a memory capable of storing a polarisation qubit was implemented by splitting the two polarisations into two spatial modes. The two modes were sent into the same ensemble, a cold cloud of caesium, achieving 17% efficiency after storage and recall

[36]. The quality of this memory was greatly improved recently by increasing the optical depth sufficiently to reduce EIT absorption, while minimising incoherent absorption [37].

In a different scheme, the clock transitions in rubidium were used to store a polarisation qubit with long storage time but low efficiency [38]. A Zeeman shift was applied to move all other transitions outside the EIT window. This increased the lifetime, but the absorption due to other transitions greatly reduced the efficiency.

A lifetime of 16 s was achieved with EIT using atoms trapped in an optical lattice [39]. A magic-valued magnetic field cancelled detunings of the memory transitions produced by the lattice trapping fields, allowing the long lifetime.

EIT was performed with the evanescent field of light travelling through a nanofibre coupled to atoms in a nearby MOT [40]. The confinement of the light in the nanofibre greatly increased the atom-light interaction, increasing the OD per atom. The coupling was improved further by trapping the atoms directly in the evanescent field of trapping lasers sent through the nanofibre [41, 42].

3.4 Echo type optical memory schemes

Echo type memories are inspired by the Hahn spin echo [43]. In this technique a $\pi/2$ pulse brings an ensemble of spins onto the equator of the Bloch sphere. The spins precess along the equator according to their individual resonant frequency. A π pulse flips the spins to the other side of the equator. Spins of higher frequency than the central frequency have advanced and so are placed behind the other spins in the mirroring operation, while spins of lower frequency are placed ahead. The higher frequency spins catch up to the lower ones and the original $\pi/2$ pulse is re-emitted. A photon echo in a ruby crystal was demonstrated in 1964 [44].

The Hahn photon echo is not a useful quantum memory on its own. A $\pi/2$ pulse for a large ensemble contains a large number of photons, and is not practical to produce as a pure quantum state. The requirement to take the ensemble through the excited state during the π pulse adds a large amount of spontaneous emission. A number of techniques that rely on the controlled rephasing of an ensemble been developed into quantum memories and are classed as photon echo type memories.

3.4.1 Atomic frequency comb

The atomic frequency comb memory (AFC) is implemented by imprinting a comb in frequency space in the absorption profile of the ensemble. This generally takes advantage of the inhomogeneously broadened absorption lines in solid state platforms. A train of optical pulses is used to write the comb by moving emitters into the ground state of the memory scheme. The time between pulses determines the frequency spacing, which also determines the delay time of the memory.

An absorbed probe is re-emitted with a certain delay $1/\Delta$ for a frequency comb spacing of Δ. The discrete frequencies in the comb rephase after a pre-determined time. The absorption only covers narrow frequency bands spaced over the larger memory bandwidth, but the whole bandwidth of an input can be absorbed and re-emitted. The finer frequency structure of the memory cannot be completely measured during the short input time, and indeed the finer structure causes the re-emission after set delay time. Memory efficiency is limited to 54% for recall in the forward direction due to the re-absorption of the light as it travels through the memory during recall [45].

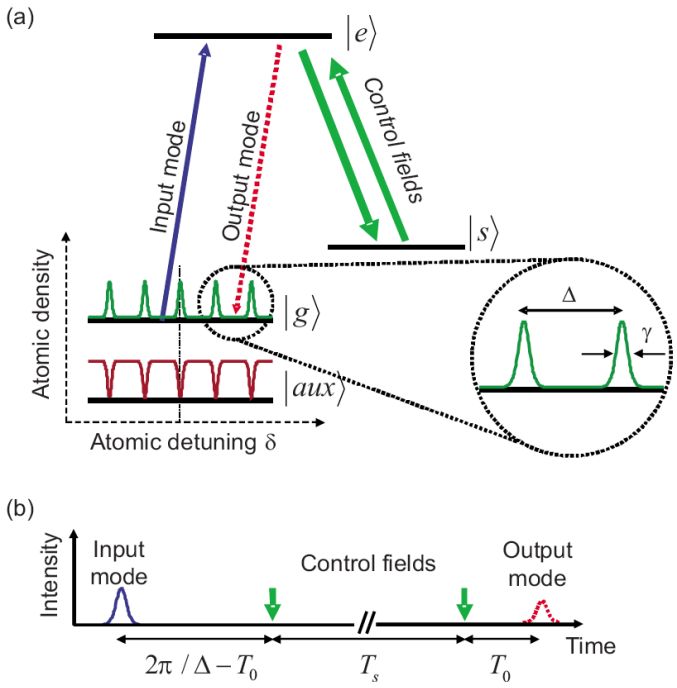

Fig. 3.1. Concept for AFC memory. a) Level structure showing absorption comb in ground state $|g\rangle$. The frequency increment between teeth Δ determines the echo interval and the width of the teeth γ the dephasing time. The control fields move the excitation $|e\rangle$ to and from the shelving state $|s\rangle$. $|aux\rangle$ provides the original broadened ensemble from which the comb is prepared. b) Timing scheme showing the delay line type storage, and the ability to recall at an arbitrary time by moving to and from the shelving state. Image taken from [45].

To create an on-demand memory, the excited state coherence can be transferred to a shelving level by use of an optical control field. The coherence can then be moved back when the output is needed. Light can be recalled in the backward direction by proper phase matching of the control fields used to transfer the spinwave to and from the shelving state. The maximum efficiency for backward recall is only limited by the optical depth. The AFC concept including on-demand memory is illustrated in Fig. 3.1.

AFC memories can be highly multi-mode due to the ability to prepare large comb structures in high optical depth materials such as rare earth doped crystals. Separate demonstrations have shown storage above the no-cloning limit [16, 46, 47], and storage of 100 temporal modes [48].

AFC can be combined with a cavity to improve the efficiency, and to exceed the 54% efficiency limit without the requirement for control fields and backward recall. A 56% efficiency was found with a weakly absorbing sample in an impedance matched cavity with a single coupling mirror [49]. In a similar setup, a control field was added to allow adjustable delay with comparable efficiency [50]. Recently, AFC was performed in a micrometer sized cavity memory in a waveguide [51]. The efficiency was only 4%, but a novel control mechanism was demonstrated: The AC-Stark shift from a strong pulse applied before recall compressed the frequency comb, allowing a tunable delay of the echo.

The coherence-preserving properties of an AFC memory were tested by simultane-

ously retrieving two pulses stored in the memory [52], with an interference visibility between the two pulses of 82%.

A spin echo technique is used in combination with an AFC memory as dynamical decoupling to improve the lifetime of the stored spinwave [53]. This kind of decoupling is used [16] with the addition of a detuning that cancels the first order Zeeman splitting, in a cryogenic solid state ensemble and yielding a T_2 of over 6 hours.

Progress toward storing telecom wavelengths has recently been made with AFC memories, which have recently been used with erbium-doped fibres to store telecom wavelength photons [54, 55, 56].

3.4.2 Controlled reversible inhomogeneous broadening

A memory based on controlled reversible inhomogeneous broadening (CRIB) also uses an inhomogeneously broadened ensemble. The continuum of frequencies means the ensemble will not periodically rephase as it does with AFC. The broadening must instead be reversed, so the inhomogeneous broadening must be entirely generated and controllable by some external field. Once the broadening is reversed, the emitters re-phase and re-emit the input in a manner similar to a photon echo. This scheme was proposed in the context of Doppler broadening [57], where counterpropagating lasers are used to apply π pulses on moving atoms that have absorbed an input photon. The inversion of the Doppler shift swaps the output direction and provides the rephasing.

CRIB memories are generally implemented in solid state. Inhomogeneous broadening arises from the sensitivity of the transition frequency to an externally applied electric field to the local strain environment. Switching the sign of the field then reverses the broadening. A CRIB memory at telecom wavelengths was implemented in erbium doped crystal [58] and the noise properties were studied [59].

Revival of silenced echo

The revival of silenced echo (ROSE) memory avoids some of the problems with the photon echo by re-emitting the input with only a small excited state population. After the input pulse is absorbed, a π pulse is used to induce a rephasing of the ensemble. However, during the first rephasing the recall of the input pulse is prevented. A surrounding cavity can be coupled and decoupled to the output [60], an applied broadening field can dramatically reduce the optical depth [61], or the phase-matching condition for the output pulse can be broken [62, 63]. A second π pulse then induces a second rephasing with the ensemble back near the ground state. The conditions for re-emission can be restored on the second rephasing to revive the echo.

The simple preparation of the ensemble allows for high optical depths and highly multimode memories, while the temporal separation between the control pulses and the output allows for low noise. The multiple rephasings increase the control pulse precision necessary for exact rephasing.

A ROSE memory was implemented in an erbium-doped crystal at telecom wavelength, with efficiency about 40% [64]. The maximum forward recall efficiency of 54% applied to this scheme, suggesting that higher efficiency could be achieved in a backward recall or cavity geometry.

3.4.3 Rephased amplified spontaneous emission

The rephased amplified spontaneous emission (RASE) memory [65] is a recently demonstrated echo-type memory. In this scheme, π pulses are used to move the entire ensemble between four different levels. An ensemble is first moved to an excited state and undergoes spontaneous emission. This spontaneous emission is amplified into the memory output mode due to the collective enhancement of the ensemble. As with other spinwave schemes, a coherent superposition of emitters in ground and excited state is generated. The emitters in the excited state are moved back down to a separate ground state and the population generated by the emission is moved to an excited state by two different frequency π pulses. The energy levels are reversed and so the incoherent broadening is also reversed, causing the coherence to rephase. The original spontaneous emission is then re-emitted as a time-reversed echo. This sequence is illustrated in Fig. 3.2. The output is temporally separated from the control pulses, reducing noise problems.

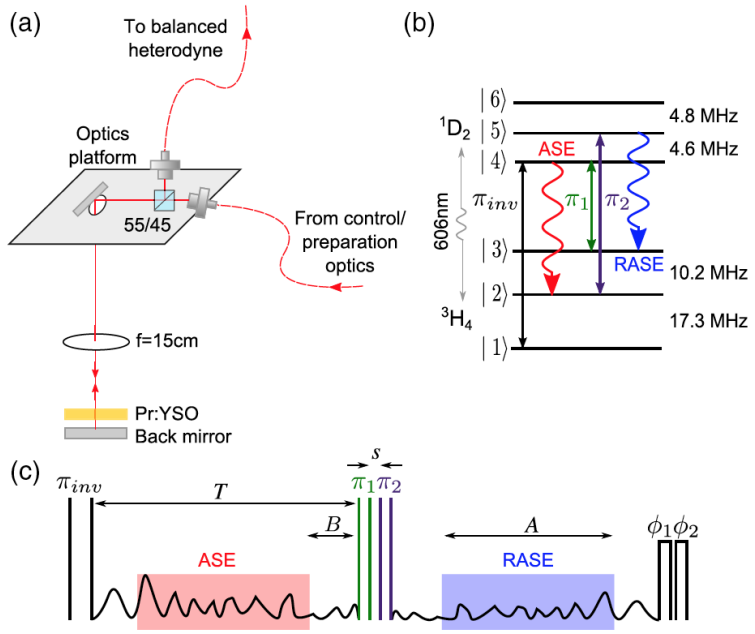

Fig. 3.2. Scheme for the RASE memory. (a) Optical setup. (b) Level scheme including the three π pulses used to move the ensemble between ground and excited states. (c) Timing of the emission with respect to the two rephasing π pulses. Figure taken from [65].

3.4.4 Gradient echo memory and Raman gradient echo memory

A gradient echo memory (GEM) is a special type of CRIB memory where the inhomogeneous broadening is purely spatially dependent [66, 34]. The broadening is still produced by an external field that shifts the resonant frequency of the ensemble. The field varies along the longitudinal axis of the ensemble, producing a gradient in the absorption frequency. The whole ensemble must have an equal sensitivity to the external field, unlike ensembles used for CRIB.

The spatial dependence of the transition frequency means that different frequency components of the input are absorbed at different positions along the memory. The emitters dephase due to the broadening in the ground state splitting and are rephased by reversing the gradient in the external field. For a monotonic gradient, re-absorption of

the recalled pulse is not an issue as the re-emitted light does not cross over atoms resonant to that frequency. The spatially dependent transition frequency and resulting spinwave is illustrated in Fig. 3.3.

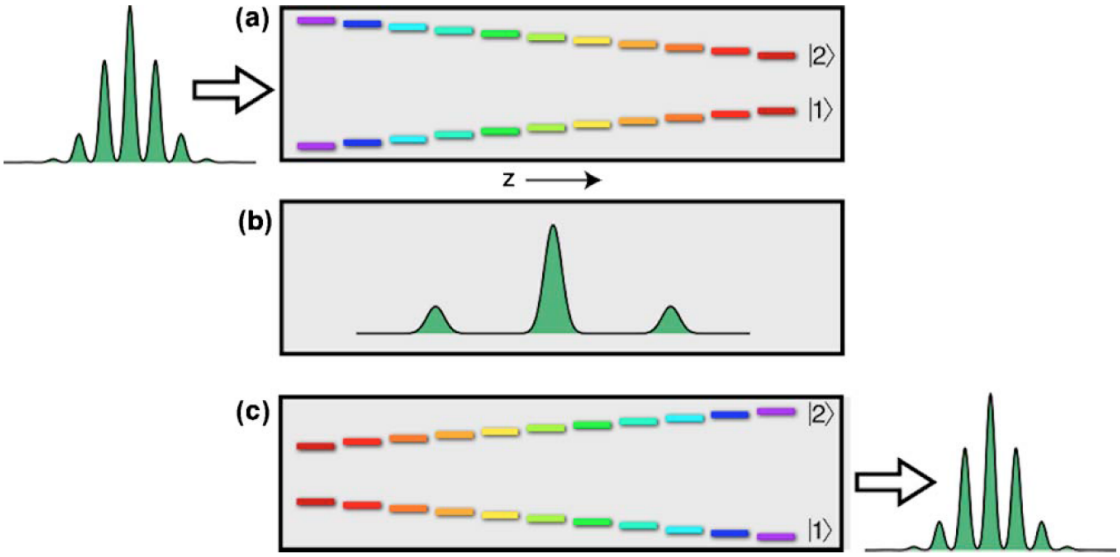

Fig. 3.3. Gradient echo memory scheme. (a) A pulse enters the memory and different frequency components are absorbed along the length z. (b) The spinwave is the Fourier transform of the input pulse. The spinwave dephases until the gradient is switched off or flipped. (c) The gradient is flipped, the spinwave rephases, and the pulse is re-emitted. Figure from [67].

GEM was demonstrated in solid state praseodymium-doped Y_2SiO_5 [68], with shifting of the levels by external electric field gradient. In 2-level GEM, the input is re-emitted upon rephasing. The spinwave can be dephased by applying a standing wave optical field, creating a varying AC-Stark shift and preventing re-emission [69]. This allows additional control over when specific temporal modes are retrieved.

Raman gradient echo memory

Raman gradient echo memory (Λ-GEM) extends GEM by coupling the input via a Raman transition to a long-lived spin coherence [70]. The Raman transition to a second ground state is completed by applying a classical laser field on two-photon resonance with the transition. This increases the lifetime of the memory and allows greater control over the re-emission of the stored light. The frequency gradient does not have to be generated by electric or magnetic fields. The AC-Stark shift from an additional laser beam with an intensity gradient along the memory has been used to create the frequency gradients [71].

Recalled pulses from GEM do not have 100% overlap with the input pulse. In short storage limit, the temporal envelope and the frequency of the pulse is distorted. This is because the light experiences dispersion as it travels through the memory toward the region of two-photon resonance. When the gradient is flipped the frequencies of the output are not linearly mapped from the input frequencies, distorting the output pulse. This can be minimised by leaving the control field and gradient on for some time after the input pulse is sent. The effect can also be cancelled by detuning the Raman transition to the opposite side of the excited state and recalling in the backward direction [72].

By treating GEM as a series of cavities, analytic functions were found that map the input pulse to the spinwave and back to the output [73]. The pulse distortion could then be studied in further detail, allowing a quantitative analysis of the fidelity of the memory.

The multiple transitions in real atoms can be used to create a more versatile GEM. For example, it is possible to store two frequency separated pulses on different transitions in rubidium [74].

The spatial multimodality in GEM was demonstrated by storing images in a warm vapour cell [75]. The combination of spatial and temporal multimodality was also demonstrated by storing and retrieving two different images sequentially [76]. Different transverse areas of the memory can be illuminated with control fields, allowing selective writing and reading of different spatial modes [77].

3.4.5 Simulated gradient memories

A simulated gradient memory creates the effect of an atomic frequency gradient by a continual change in the spatial phase-matching condition. The spatial dependence of the change in the phase-matching produces a spatial dependence of the frequency and an effective gradient in the absorption frequency of the medium. The continuous change also means that the spinwave momentum to which the probe couples changes continuously with time, allowing the pulse sequencing and other manipulations possible in GEM. Control field angular scanning is shown in Fig. 3.4.

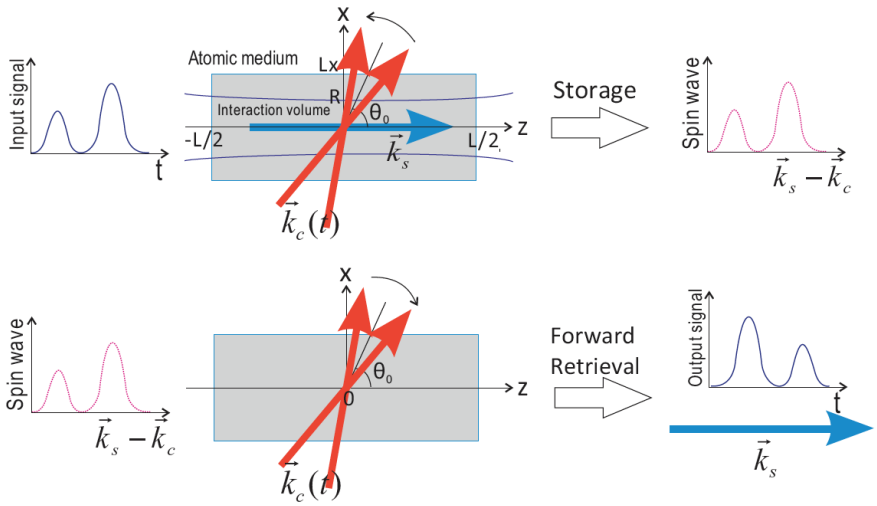

Fig. 3.4. The spinwave k-vector that the incoming signal couples to is continuously modified by scanning the control field angle θ. The position dependent movement of the phase fronts of the control field sets up a spatial frequency gradient in the absorption profile, as in a gradient echo memory. Figure from [78]

An implementation of the simulated gradient was proposed by scanning the angle of the control field [78], or changing the dispersion of the control field [79] or the input field [80] travelling through the memory. Angular scanning memories are particularly sensitive to the transverse extent of the stored mode as the spinwave wavelength can be quite short even at small angles. Changes in the transverse projection of the spinwave momentum affect the efficiency of the memory.

3.5 Raman memory

If the timing of the input state is well known, the requirement to store from a larger bandwidth is relaxed and the memory response can be tailored to store the specific temporal mode. For example, in the absorptive Raman regime the strength of the coupling and so the amount of light absorbed or re-emitted from the memory can be modified by changing the intensity of the control field. This method was proposed as a quantum memory in 2000 [81]. An analytic approach was used by Gorshkov et al. to calculate the optimal efficiency of such a memory [7]. The propagation of the pulse within the memory can be calculated as that of a pulse in a narrow linewidth absorbing medium [82].

Relying entirely on optical fields to control the coupling can allow for very fast memories, with bandwidths exceeding 1 GHz stored in caesium vapour [83]. Efficiency from the same group was improved to over 30% with experimental improvements including increased control power [84]. The setup was also modified for dual-rail storage of polarisation [85]. Photons generated by SPDC were stored in warm atomic vapour, and the effect of added noise from four wave mixing was explored [86].

Raman memories have been implemented in other systems including cold atoms [37] and diamond, where the ground states are separated by phonon creation, effectively storing the photons as phonons [87]. A hollow core photonic fibre loaded with caesium atoms was used to store single photons with an efficiency of 27% and a SNR of 2.6:1 [88].

3.6 DLCZ

The DLCZ protocol, named after Duan, Lukin, Cirac, Zoller, who proposed it [89], is an output only memory; the input is created within the memory. An idler photon is generated from an ensemble probabilistically by a spontaneous Raman transition while illuminating with a write laser field. A spinwave is generated in the memory, and by illuminating again with a read laser field, the spinwave will be read out, emitting a signal photon. This process is illustrated in Fig. 3.5.

The efficiency for generating heralded single photons is limited by the requirement to only generate a single write photon. This limits the illumination with the write laser field and thus the probability of photon generation. However, recall efficiencies once a photon is generated can be very high.

A DLCZ scheme was demonstrated that took advantage of the spatially multimodal emission of write and read photons [91]. Multiple emitted write photons in spatially separated modes heralded the creation of multiple read photons. These read photons were recalled into predictable spatial modes that could potentially be routed into a circuit requiring multiple input photons.

The lifetime of a DLCZ memory was extended using a scheme that could also be extended to other spinwave-based memories. Two additional laser fields were used to transform the spinwave generated in a cold atom DLCZ memory to a coherence on a magnetically-insensitive transition while also minimising the spinwave momentum [92]. This extended the memory lifetime until it was only limited by the overall expansion of the atom cloud.

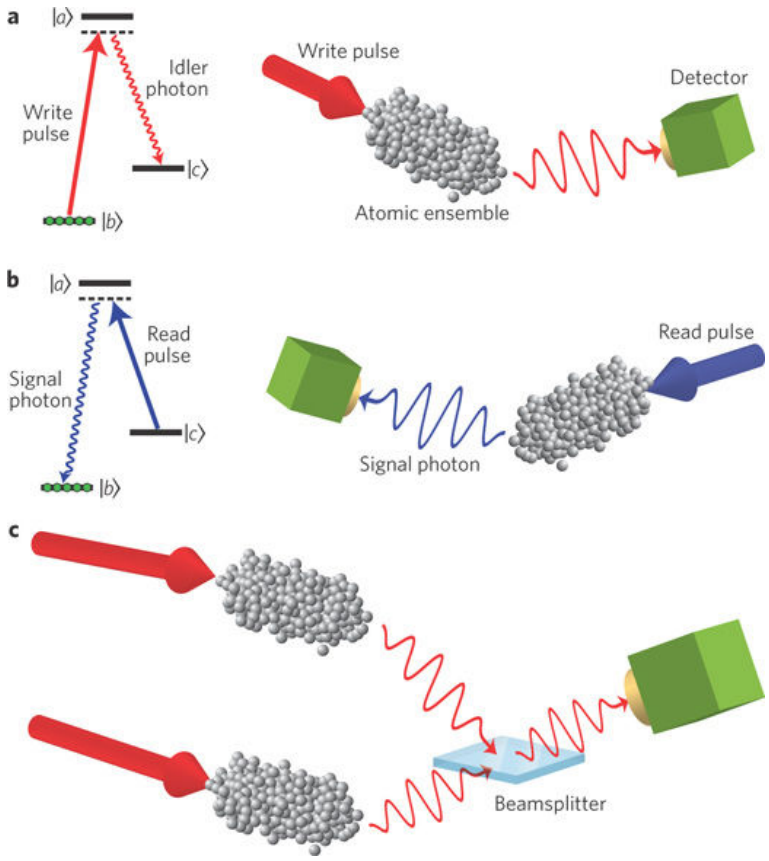

Fig. 3.5. In **a**, a spinwave is prepared by the probabilistic scattering of an idler photon from the write pulse. In **b**, a signal photon is emitted from the ensemble by reading the spinwave with a read pulse. In **c**, two distant ensembles are entangled by combining the write photon mode at a beamsplitter. Detection of a photon projects the ensembles into an entangled state, with the spinwave in one ensemble OR the other. Figure from [90]

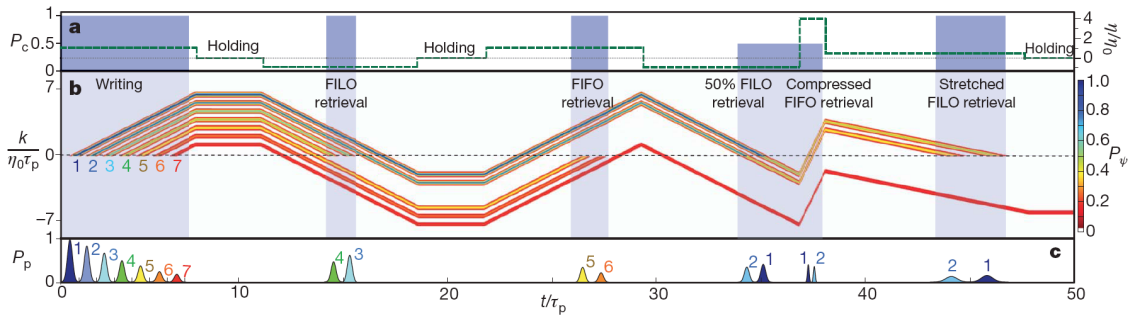

Fig. 3.6. The top row **a** shows the coupling field and gradient. **b** shows the spinwave in momentum. The time envelope of the pulse becomes the momentum envelope, which is transported according to the gradient. Light is recalled if a spinwave rephases (crosses zero momentum) while the coupling field is on. The proportion of light recalled depends on the coupling field strength, while the gradient magnitude can be modified to change the bandwidth.

3.7 Single mode memories

Many unbroadened ensemble memories are capable of being turned into single mode. By placing the memory inside a cavity, spatial variation along the length of the memory is removed, as light travels through the memory many times while being absorbed. The memory becomes more efficient for a given optical depth compared to a multi-mode memory [93] and more selective in what temporal mode is stored. In 2016 a cavity enhanced Raman memory was demonstrated with a SNR of about 6 at single photon levels and an efficiency of 9.5% [94]. A DLCZ memory was implemented in a cold atom ensemble in a cavity, with a 73% retrieval efficiency of the read photons [95]. In a similar setup, 82% efficiency with 900 ns storage was more recently achieved [96]. By completely averaging out spatial information, the lifetime of memories based on warm atoms can be dramatically increased. For example, Borregaard et al. demonstrate a single mode memory based on a warm atomic vapour in a micro-cell with a special coating that preserves the atomic spin upon collisions with the walls [97].

3.8 Coherent processing

A collection of memories and single photons has been proposed as a useful multi-photon resource [10]. However, a single memory can also be used where that memory is capable of storing many single photons. Sufficient efficiency, lifetime, and control over input and output could allow the synchronisation of many photons without the infrastructure of many memories and sources. Some memories can also be used to modify the spectral properties of their outputs, or effectively as tunable bulk optics to increase the flexibility of an optical setup.

The gradient echo memory displays this versatility. It has been used as an optical sequencer, allowing the reordering of input pulses [98]. Fig. 3.6 shows how the transport through momentum space of the spinwaves generated by each pulse allows the selective recall of each pulse dependent on the gradient and coupling field switching. For example, the pulses are recalled in reversed order after the first gradient switch (first-in-last-out, or FILO) and in input order after the second gradient switch (first-in-first-out, or FIFO).

It also shows the ability to partially recall pulses by adjusting the coupling fields strength, and to change the output bandwidth by changing the gradient magnitude.

Further demonstrations of this versatility include the proposal to use the spatial separation of frequencies in the memory to apply different gradients to different regions of the memory, allowing the separation or combination of different frequency components [67]. This was carried out by using multiple individually drivable coils to create the variable gradients [99].

Sending a pulse into a memory while a pulse is simultaneously recalled can result in an interference between two outputs. The transmission of the sent pulse interferes with the recalled pulse, while the stored part of the sent pulse interferes with the remaining part of the pulse originally in the memory. This interference is phase-dependent and closely resembles the interaction of pulses sent into the two inputs of a beamsplitter, with a reflectivity tunable by the coupling into and out of the memory. This was demonstrated in the gradient echo memory [100]. Beamsplitter operation on the output with no additional inputs was demonstrated in a Raman memory, with multiple partial recalls of the input [101].

A protocol for quantum computation only using linear optics and projective measurements, known as linear optical quantum computing was proposed [102]. It was shown that the gradient echo memory could be extended to act as an array of configurable beamsplitters, making such a scheme more practical [103]. Optical processing has also been demonstrated much earlier with photon echoes [104].

3.9 Frequency conversion

Frequency manipulation can be taken a step further by the use of nonlinear processes to convert photons to very different wavelengths. This allows the use of wavelengths compatible with the various memory platforms and wavelengths compatible with communication. Photon generation, processing, and storage can be done at the memory wavelengths, and the results converted to and from telecom wavelengths for long distance communication.

Demonstrations of conversion to the telecom band at the single photon level include the conversion of rubidium wavelength photons [105], 738nm weak pulses [106], weak pulses and photons from SPDC [107]. The single photon character of the output is precisely measured with superconducting single photon detectors [108].

One photon from a pair was converted to telecom wavelength and the other stored in an AFC memory [56]. Correlations in the detection of the photons were preserved after accounting for the storage time. These demonstrations have made the use of memories outside telecom wavelengths seem more feasible.

3.10 Summary

The field of optical quantum memories is large and rapidly developing. This review gives context for the work in the following chapters in which we characterise an optical quantum memory, use a memory as an optical resonator, and analyse the operation of a gradient echo memory with a single spatial dimension.

Chapter 4

An optical quantum memory in cold rubidium atoms

The gradient echo memory scheme has been used with great success in our group at ANU for the last decade. The scheme has been used with both warm and cold atomic ensembles to demonstrate highly efficient optical quantum memories in addition to a range of coherent optical processing schemes, as discussed in Chapter 3.

This chapter is about the implementation of an optical quantum memory in a high optical depth cold atomic ensemble using the gradient echo memory protocol. The published paper relevant to this chapter is

"**Highly efficient optical quantum memory with long coherence time in cold atoms** Y.-W. Cho, G. T. Campbell, J. L. Everett, J. Bernu, D. B. Higginbottom, M. T. Cao, J. Geng, N. P. Robins, P. K. Lam, B. C. Buchler, Optica 3, 100-107 (2016)"

The above paper includes a high efficiency demonstration, presenting the first memory to outperform the theoretical limit for a low-loss optical fibre delay line. If this does not sound particularly impressive, it is worth mentioning that it would be nearly impossible to actually reach the theoretical limit in fibre and it would not produce a good quantum memory. Long lengths of fibre are sensitive to acoustic and thermal noise, and are not on-demand memories without significant modification.

The method used to characterise the preservation of the quantum information was to extract the Wigner functions of stored and retrieved states and measuring the overlap, as well as analysing in the T-V formalism. A large number of coherent states were stored in the memory and retrieved. The outputs were measured by heterodyne detection, allowing a direct sampling of the Q function. This allowed the phase and amplitude quadratures to be measured. The size of the initial and final states, along with the variance in each quadrature, could be established with sufficient statistics, allowing reconstruction of the Wigner functions and T-V analysis.

4.1 Introduction to high efficiency GEM in cold atoms

A long lifetime, high efficiency gradient echo memory based on atoms was implemented in a cold cloud of rubidium-87 atoms. The cloud of atoms was generated using a magneto-optic trap (MOT). A MOT uses magnetic fields and lasers to both create a trapping potential and to cool the atoms. A temporal dark spot was used to compress and cool the atoms, producing an amplitude optical depth, equivalent to the d in Eq. (2.72),

on the probe transition of about 250 and a temperature of about 100 μK. A detailed account of building and characterising the MOT can be found in the thesis of Ben Sparkes [109]. More information about the exact arrangement and timing of fields can be found in the paper for this chapter [110].

Once the high optical depth ensemble of atoms was prepared, most of the atoms were pumped into the ground state $\left|5^2S_{1/2},\ F=1,\ m_F=+1\right\rangle$ using a circularly polarised resonant with the transition $\left|5^2S_{1/2},\ F=1\right\rangle \rightarrow \left|5^2P_{3/2},\ F=2\right\rangle$ along with a repump laser resonant with $\left|5^2S_{1/2},\ F=2\right\rangle \rightarrow \left|5^2P_{3/2},\ F=2\right\rangle$ to clear atoms from the $\left|5^2S_{1/2},\ F=2\right\rangle$ state. This preferentially moved atoms to the desired ground state. This allowed the use of the level scheme shown in Fig. 4.2 to implement the Raman-GEM protocol. Fig. 4.1 is a photo of the atomic cloud in the vacuum chamber with the magnetic field coils.

The approximate layout of the experimental fields is shown in Fig. 4.3. A probe field was generated from an MSquared Ti:Sapph laser. To achieve the stability required for the quantum characterisation, the laser was locked to the transition $|2\rangle \rightarrow |3\rangle$. A few mW of laser power was sent through a fibre-coupled waveguide electro-optic modulator (EOM), modulated at 6.84 GHz. A mode cleaner cavity was used to select the +6.84 GHz sideband, which was then used for the probe. The intensity time profile of the probe was shaped using an AOM.

The control field was generated from the original laser frequency. An AOM was used to detune the control field 160 MHz from resonance and to shape the intensity time profile. The EOM was then used to tune the probe frequency to two-photon resonance, using the unbroadened Raman absorption line as a guide.

The magnetic field gradient was generated by a pair of coils, coaxial with the MOT and situated equal distances along the axis from the centre of the memory. This generated a nearly linear frequency gradient when different currents were sent through each coil. The gradient could be reversed by swapping the currents. The currents were driven by operational amplifier circuits in a voltage-controlled current source configuration.

A probe pulse with a Gaussian time profile was stored while the control field was sent and a magnetic field gradient applied. The control field and gradient were switched off after storage. The pulse was later recalled by applying the reverse gradient and sending the control field. An example of the timings for these fields is shown in Fig. 4.4. In an alternative scheme for the shorter storage times, the control field was left on and the gradient was reversed between storage and retrieval without switching it off.

The efficiency was measured by comparing the integrated intensity of the input pulse with no atomic cloud with intensity of the recalled pulse as measured at an avalanche photo-diode (APD).

4.2 Optical quantum memory characterisation

For the quantum characterisation the memory output was mixed with a bright local oscillator at a beamsplitter, making a heterodyne measurement. The extent to which the quantum state of an input was preserved could be calculated after taking a large number of these measurements.

Heterodyne measurement was limited by experimental constraints. The high efficiency data was able to be taken at a larger detuning, allowing for a higher efficiency measurement of up to 87% as measured by integrating the intensity of the input pulse and the retrieved pulse. However, the longer time necessary to build up statistics for

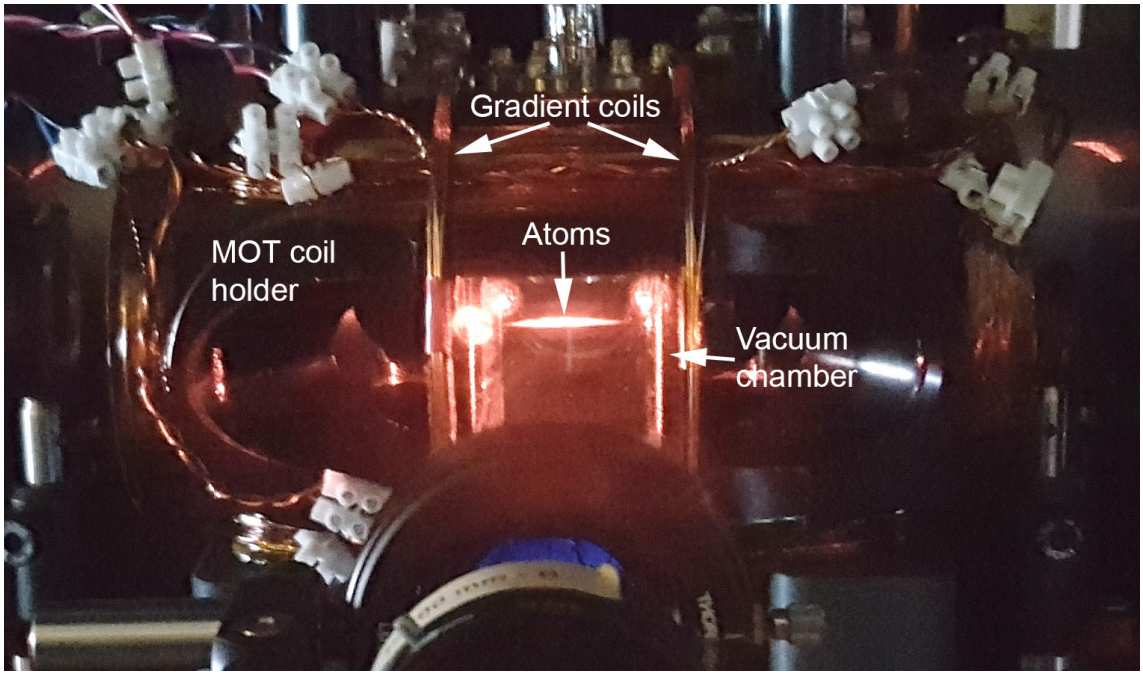

Fig. 4.1. Detail of the vacuum chamber and coils. The racetrack shaped coils and the circular end coils generate the elongated magneto-optic potential. The coils just to either side of the vacuum chamber are in anti-Helmholtz configuration and, along with the bias coils (not shown) generate the magnetic field for the gradient echo memory. The image is taken during the whole trapping cycle, and the atoms fluoresce in the trapping fields. During a memory experiment, the trapping fields are off and the atoms are much more tightly compressed.

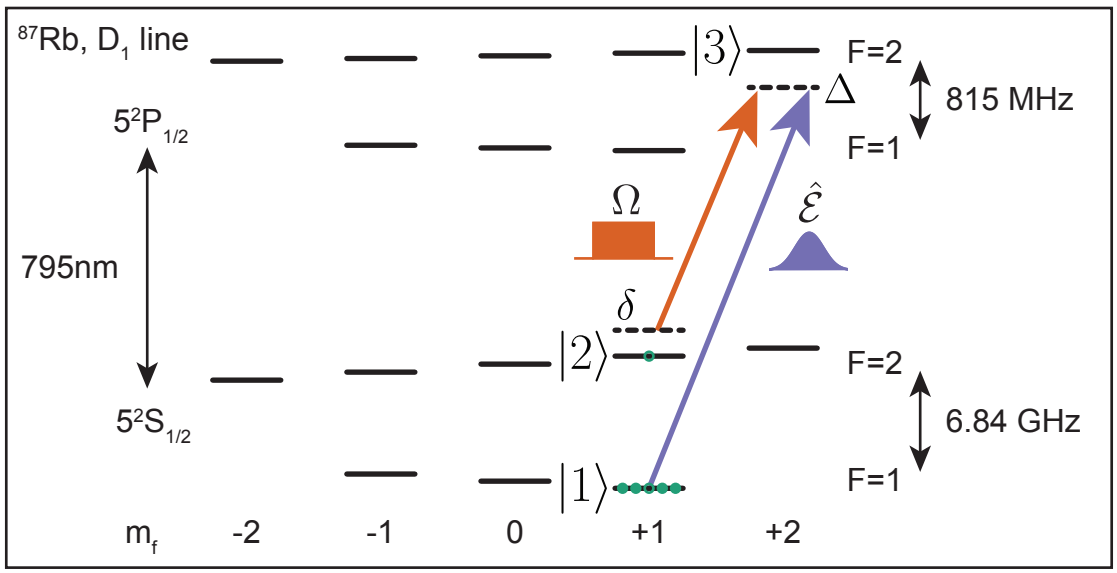

Fig. 4.2. Level scheme for the gradient echo memory experiment. Atomic energy levels taken from Steck alkali data [111].

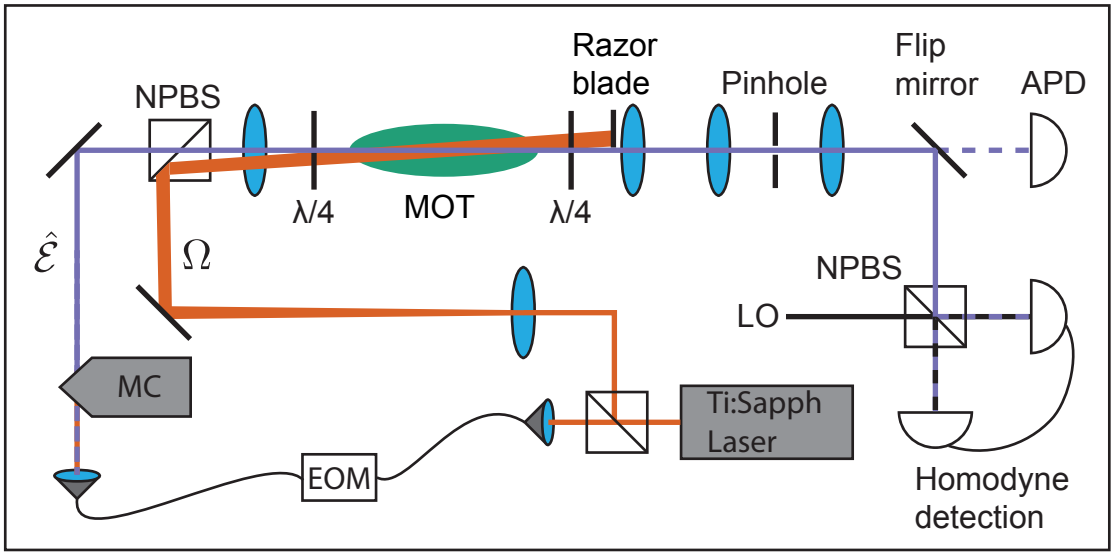

Fig. 4.3. Setup for cold GEM. The probe beam is focussed at the centre of the MOT, then the control field is removed with a razor blade. The control field is filtered further by refocussing the probe through a pinhole. The transmitted probe is measured either with an APD or homodyne detection.

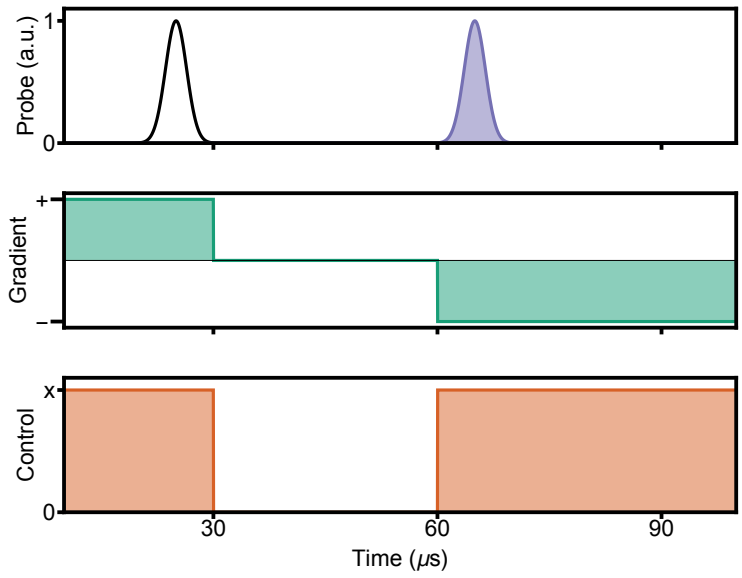

Fig. 4.4. Example control field and magnetic gradient timings for the gradient echo memory experiment. A pulse (black) is sent and the control field and gradient switched off. The echo (blue) is retrieved after some storage time by switching the reverse gradient on along with the control field.

tomographic measurement required a more stable memory. This was achieved by locking the probe laser to a specific rubidium transition, reducing the possible detuning to 200 MHz due to the bandwidth of the control AOMs, resulting in a less efficient memory. While the efficiencies for this data are not as high as for the high efficiency results, a quantum memory was still established out to times of 125 μs. This outperformed an ideal optical fibre delay, which has a maximum of 100 μs while remaining in the no-cloning regime.

4.2.1 Results and analysis

Example raw heterodyne measurements are shown in Fig. 4.5. The phase of the measurements was referenced to the heterodyne measurement of a bright probe pulse sent through the memory prior to each storage and without a control field. The quadratures were extracted by demodulating the signal with sine and cosine functions with a Gaussian envelope. The phase offset of these functions was determined by the phase of the reference pulse. The frequency chirp due to the GEM recall process was accounted for in the envelope function. For example in the 10μs storage data a chirp of -0.32 MHz/μs was used. The measured quadratures for three different input states are shown in Fig. 4.6.

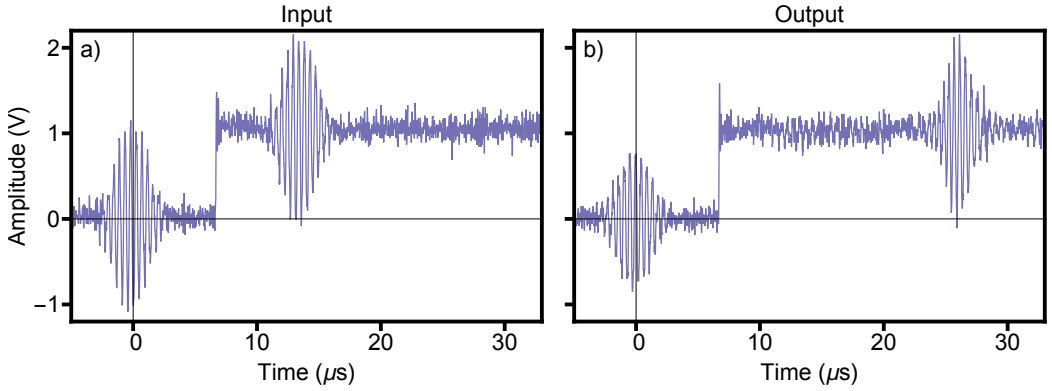

Fig. 4.5. Example heterodyne measurement of a very bright state; a) the unstored input state and b) the recalled state. The phase of the modulation is extracted and compared to the reference pulse centred at time $t = 0$, and the amplitude is extracted by fitting a Gaussian envelope. The jump in amplitude is due to control field switching on and leaking through to the detector.

There are many benchmarks against which to measure the performance of a quantum process. We chose to use the $T-V$ formalism to measure the performance of the memory. This allowed us to compare the performance against that of an ideal linear process. The $T-V$ measure is an attractive benchmark because it is state independent. We can then be confident that the memory will preserve any state that fits within the memory bandwidth, and not just the state we have chosen to use for measurements. This should be taken with the caveat that specific uses of the memory may still be more sensitive to added noise such as from the control field.

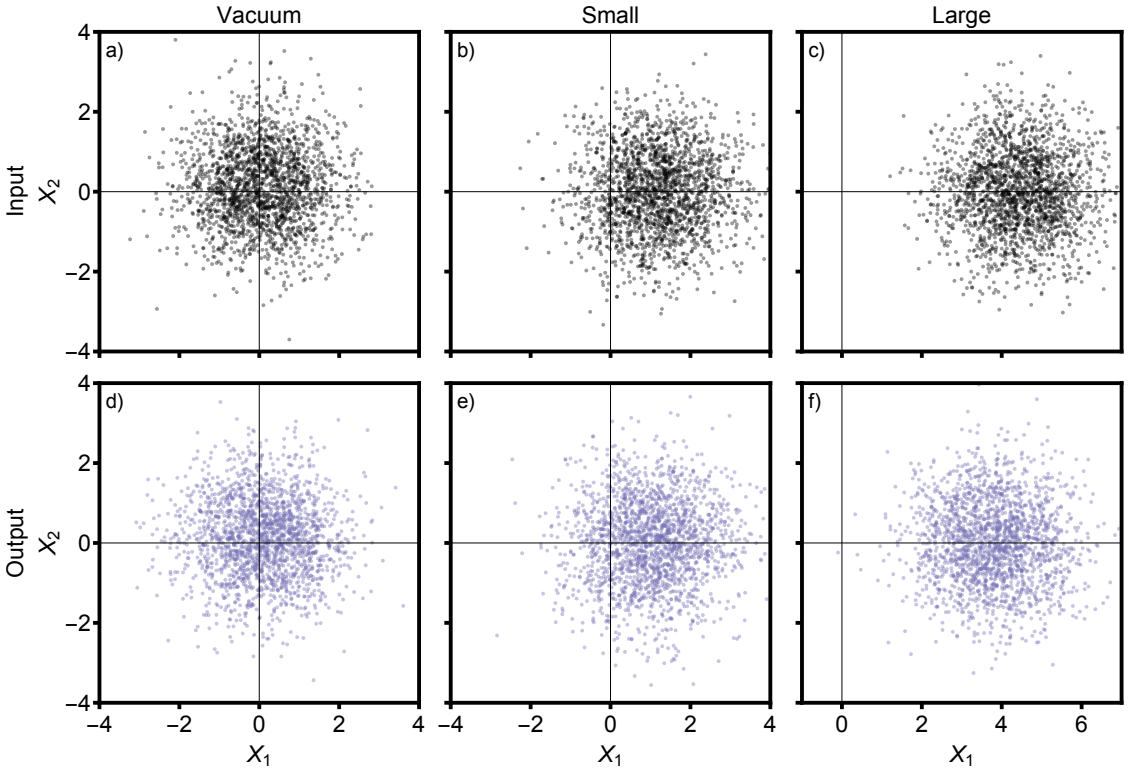

Fig. 4.6. Quadratures extracted from 10 μs storage data in units of shot noise for the various sized coherent states sent and recalled from the memory. Top row are the input states and bottom row are the outputs. Each point corresponds to a single sent pulse, with the value in each quadrature corresponding to the overlap with the corresponding sinusoidal envelope. Each set is brought to the X_2 axis by a rotation based on the average angle of the set.

Wigner functions

Wigner functions were also extracted from the quadrature data by deconvolving the data with a Gaussian kernel. Fig. 4.7 plots the Wigner functions extracted from the heterodyne data. These show in particular the preservation of the phase of the coherent state, as the recalled state stays close to the $X_2 = 0$ axis.

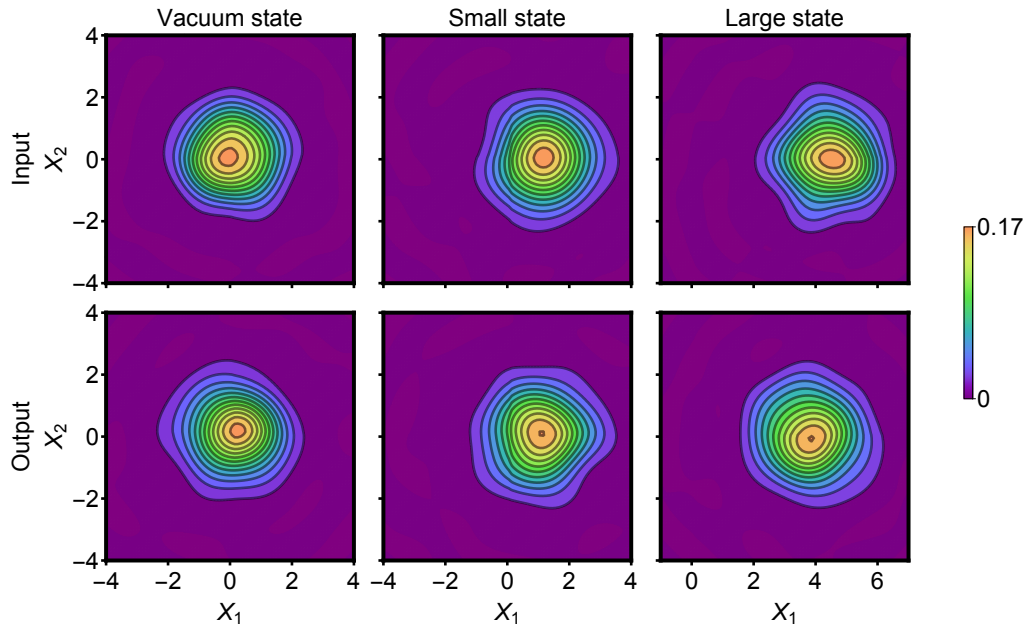

Fig. 4.7. Wigner functions $W(X_1, X_2)$ extracted from the quadrature measurements for the 10 µs storage data.

T-V analysis

The $T - V$ formalism was developed for characterising teleportation [112], and has been adapted for characterising memories in [34, 113]. However those methods do not account for the efficiency of the detection and the noise must be scaled accordingly. For example, if the detection was very inefficient, the noise added by the memory would also be very small and the memory would seem less noisy. The variances are corrected by transforming the measured quadrature variances $V_q^{\pm}$ in the following way:

$$V_{in/out}^{\pm} = \frac{V_q^{\pm} + \eta_D - 1}{\eta_D} \tag{4.1}$$

where η_D is the detection efficiency. The conditional variance for each quadrature $V_{CV}^{\pm}$ is given by

$$V_{CV}^{\pm} = V_{out}^{\pm} - \frac{\left|\langle \delta X_{in}^{\pm} \delta X_{out}^{\pm} \rangle\right|^2}{V_{in}^{\pm}} \tag{4.2}$$

with $\delta X^{\pm}$ as the measured quadrature variances.

For a storage efficiency η_S, calculated from the photon numbers of the input and

output states, the $T - V$ parameters are

$$V = 1 - \eta_S + \sqrt{V_{CV}^{+} V_{CV}^{-}} \tag{4.3}$$

$$T = \frac{2\eta_S}{1 + \sqrt{V_{CV}^{+} V_{CV}^{-}}}, \tag{4.4}$$

The V parameter can be understood as the vacuum noise added due to the inefficiency of the memory $(1 - \eta_S)$, as well as the noise added by the memory itself. The T parameter is similar to a signal-to-noise ratio.

The above method of averaging the conditional variances relies on the assumption that the storage and recall processes are not dependent on the phase of the input. This assumption was justified by analysis of the measured quadratures. The parameters for several datasets are shown in Fig. 4.8. The error bars for the 125 μs dataset are smaller due to a brighter local oscillator and a larger number of samples (4000 compared to the 2000 each for the other datasets).

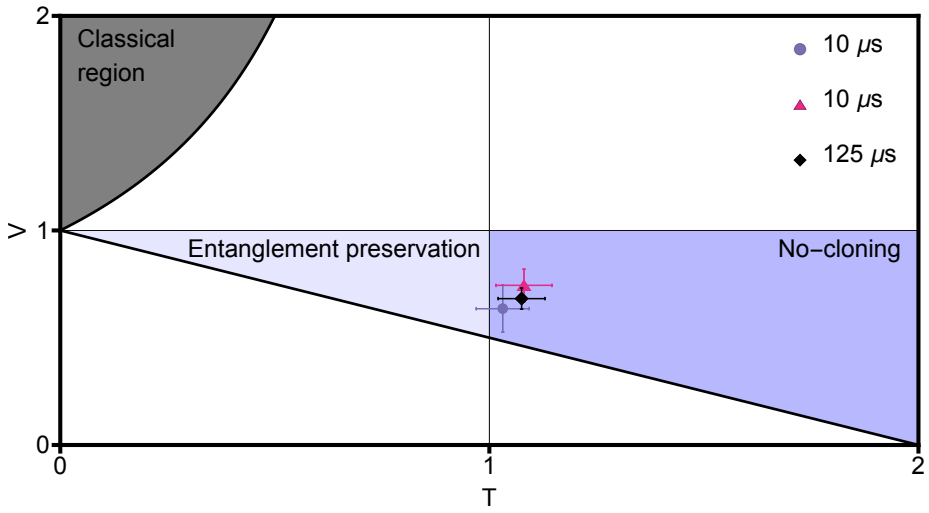

Fig. 4.8. T-V diagram of the memory process for different storage times. Two different datasets for 10 μs storage are shown. The data for 125 μs of storage place the T-V point more than one standard deviation inside the no-cloning condition (blue shaded region). The lower black line marks the boundary where there is no added noise and the T-V parameters correspond to linear loss. The upper, curved black line marks the limit of how accurately the state could be recreated after measurement.

The T-V parameters place the data inside the no-cloning region. This means that the output from the memory is the best possible copy of the input state. This is an important bar to reach for the security of quantum communication protocols based on the memory, and we take satisfaction of this condition to be necessary to call the memory 'quantum'.

We were able to show a quantum memory even out to storage times of 125 μs. The heterodyne data shown for 10 μs shows 73% efficiency while for 125 μs the efficiency was 53% and the highest efficiency measured on the APD data was 87%. This suggests that the memory could easily be pushed further into the no-cloning regime by operating at a larger detuning. As above, the requirement for frequency stability restricted the possible detunings for the heterodyne measurement data.

The main source of noise the memory produced on the state were shot to shot variations as the atomic cloud was released and then cooled and compressed again for each storage. This cycle produced fluctuations in the magnetic field, which can affect the phase with which the pulse is recalled. The variation in optical depth also affects the phase, as this determines the strength of the coupling into the memory. Shot to shot variations in the phase manifest as a increased spread in the distribution of the heterodyne data in the phase quadrature. The consequent increase in the conditional variances $V_{CV}^{\pm}$ moves the data away from the linear loss line in Fig. 4.8.

The dispersion experienced by the probe pulse moving through the memory would also affect the phase. However, this would likely affect the reference pulse by a similar amount, reducing the apparent fluctuation of the phase. The dispersion is easily measurable without destroying the MOT, so it is not unfair to neglect this variation.

4.3 Conclusion

In this chapter we demonstrated an optical quantum memory in cold rubidium-87 atoms. We used heterodyne measurement to demonstrate that the quantum states of weak coherent pulses stored in the memory were sufficiently preserved to satisfy the no-cloning condition. The conditional variance can be reduced by reducing shot-to-shot variations in the atomic ensemble. Improving the efficiency of the memory further will require experimental improvements such as increasing the optical depth and increasing the one-photon detuning. Increasing the detuning in particular will require more isolation of the probe from the control field, with frequency filtering from etalons or warm rubidium vapour. A low control field transmission will also be necessary for improving the signal-to-noise ratio when storing single photons.

Chapter 5

A spinwave resonator

> Someone once told me, time is a flat circle
>
> Rust Cole, *True Detective*

In Chapter 3 we covered some of the uses of multi-mode memories for spectral processing of optical inputs. In particular, the partial absorption and re-emission of inputs in a multi-mode memory allows for the implementation of beamsplitter-like operations.

In this chapter we extend the beamsplitter concept of Campbell et al. [100] to construct the analogue of a pulsed optical resonator. The stored memory mode replaces the circulating optical field in an optical resonator built with mirrors. In the memory-based resonator, the portion of the light pulses transmitted through the memory interfere with the simultaneously recalled mode, while the portion coupling into the memory interferes with the stored mode. Fig. 5.1 illustrates the conceptual link between a resonator based on a circulating optical mode and the spinwave resonator. The repeated interference between a series of input pulses and a circulating mode is common to both, though the mode is optical in the case of the optical resonator, and is a spinwave converted to an optical field in the case of the spinwave resonator.

The relevant published paper for this chapter is

"**A mirrorless spinwave resonator** O. Pinel, J. Everett, M. Hosseini, G. T. Campbell, B. C. Buchler, P. K. Lam, Scientific Reports 5, 17633 (2015)"

The coupling into and out of the stored mode can be adjusted with control field intensity, allowing for dynamic tuning of the resonator properties. The repetition rate of the resonator is also adjustable by the switching time of the magnetic field gradient, and the bandwidth and central frequency of the interaction depends on the magnitude of the gradient and the constant offset of the magnetic field.

5.1 Method

The spinwave resonator is implemented in the gradient echo memory protocol in a warm vapour of rubidium-87. A series of pulses is sent into the memory, with the recall from the memory synchronised with the input of the next pulse.

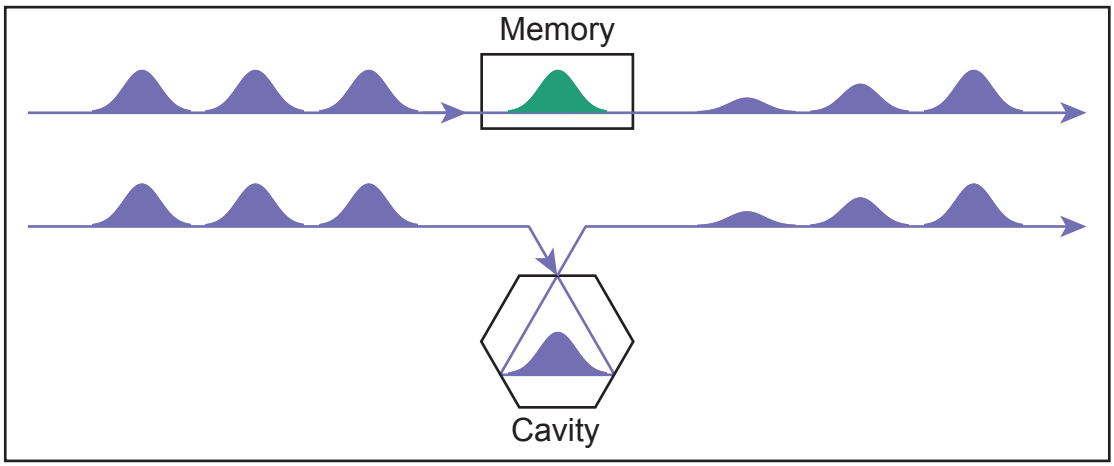

Fig. 5.1. Concept for an optical resonator based on gradient echo memory. The memory acts as a cavity with a tunable coupling. A cavity with a single coupling mirror is comparable to the memory if the round trip time is longer than the input pulses. The blue Gaussians are the input and output optical pulses and the pulse circulating in the cavity. The green Gaussian represents the spinwave in the memory, which takes the place of the circulating optical pulse.

5.1.1 Experimental setup

The experimental setup is illustrated in Fig. 5.2, consisting of a similar physical setup to that used in previous warm vapour GEM experiments, eg. Hosseini et al. (2009) [98].

Warm rubidium vapour

The vapour cell was 20 cm long with a diameter of 2.5 cm and contained a small drop of isotopically enhanced rubidium-87 and 0.5 Torr of Kr buffer gas. The cell was heated to a temperature of about 75 °C which generated about 1.25×10^{12} atoms/cm^3 [111]. The temperature was chosen as a compromise between optical depth which improves the performance of the memory, and Doppler broadening and atomic diffusion which reduce the memory performance. The cell was heated with a resistive non-magnetic wire, with the current switched off 2 ms before each experiment to allow magnetic fields generated by the heating current to die down. The heater was switched back on immediately after the experiment and so spend 90% of the time switched on.

Magnetic fields

To eliminate background magnetic fields and ensure the magnetic field points along the optical axis, the vapour cell was housed in a shield of two layers of μ-metal. This high permeability metal diverts the external background magnetic field through itself, greatly reducing the magnetic field inside the shield.

A coil with a winding of constant pitch was used to apply a magnetic field of about 6 G at the long axis of the cell. The magnetic field gradient was produced by a pair of coils, each with varying pitch. Each coil mirrored the other, so that the sign of the gradient could be flipped by switching one coil off and the other on. This method of gradient flipping avoided changing the atomic frequency offset, ensuring the carrier frequency of

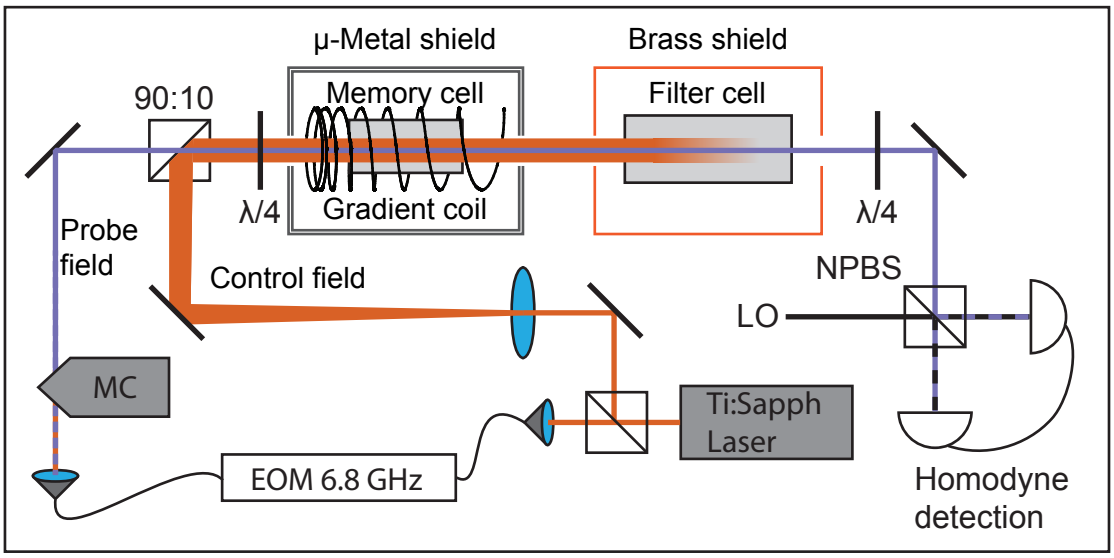

Fig. 5.2. Experimental setup for the spinwave resonator. The electro-optic modulator (EOM) phase modulates the original laser frequency. The +6.8 GHz sideband is separated at the mode-cleaner ring cavity (MC) and used as a probe field. The probe field and control field are combined and sent through the memory. The control field is filtered out in a cell containing rubidium-85. The remaining probe is combined with a local oscillator also derived after the EOM. Two gradient coils are used but only one is shown, the other has a pitch increasing in the opposite direction. A coil with constant pitch, not shown, is also used to apply a constant magnetic field.

the recalled pulse matched that of the sent pulse. The magnetic field gradient produced by the gradient coils was 5 G/m, producing a broadening of 1.4 MHz of the Raman absorption line for a 20 cm cell. The coils were nested with the cell at the centre and PVC tubes maintaining the coil spacings and the coaxial alignment.

Laser frequencies and pulses

The laser fields were generated from an MSquared Ti:Sapphire laser tuned near the rubidium D1 line at 795 nm. The laser was detuned by 2 GHz above the $F = 2$ to $F' = 2$ transition by monitoring the fluorescence of the laser in a cell containing rubidium of natural isotopic ratio.

The Ti:Sapphire laser was not locked to a particular atomic resonance so the one-photon detuning is not precisely known. Instead, the laser was locked to an internal reference cavity. The frequency stability of the laser was then sufficient to produce homogeneous data over the time period of the entire data collection.

The probe was generated by sending a few mW of laser power through a fibre-coupled waveguide EOM (EOSPACE), modulating the phase of the light at approximately the ground-state hyperfine splitting of 6.8 GHz. The +6.8 GHz sideband was selected using a mode cleaner triangular cavity. Probe pulses were generated at an acousto-optic modulator. A local oscillator used for heterodyne measurement of the probe was also split off at this point and detuned 2 MHz from the probe. The relative detunings of the probe and control fields from the coupled transitions are shown in Fig. 5.3.

The probe beam was sent through a cell of isotopically enhanced rubidium-85 heated

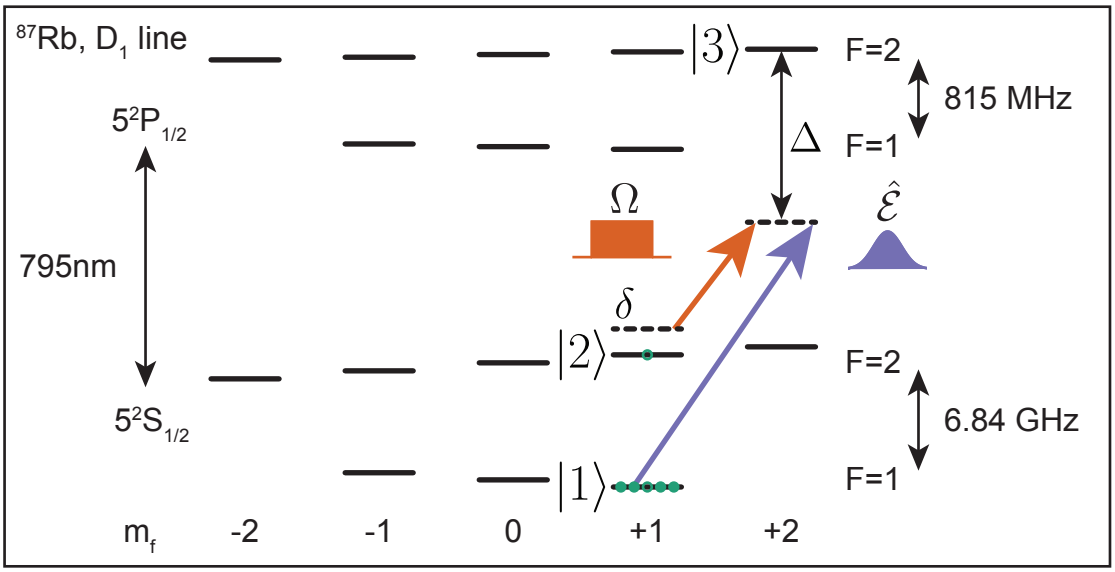

Fig. 5.3. Energy level scheme for rubidium-87 with relative frequencies and polarisations of experimental optical fields. Solid green circles represent the initially occupied level, with empty circles representing the atomic coherence. Atomic energy levels taken from Steck alkali data [111].

to 150 °C. The control field was close to the resonance between $F = 1$ to $F' = 1$ in rubidium-85 and was strongly absorbed. The probe field was far from any resonance and was only moderately absorbed in the filtering cell.

Atom preparation

The control field was left on between experiments at a higher intensity to pump atoms into the $F = 1$ state, with preferential pumping due to the control field polarisation into the $m_f = +1$ Zeeman level and the splitting of the Zeeman levels. This prepared an effective three-level atomic ensemble. Atoms in other ground levels still caused some absorption and dispersion of the probe but did not contribute to any coherent evolution.

Measurement

The probe was measured after the filtering stage by heterodyne measurement. Heterodyning helped to prevent the remaining control field background from affecting the measurement. Only the signal of the interference between the local oscillator and the probe within the detection bandwidth was kept. The lower bandwidth signal corresponding to any intensity fluctuations of the transmitted control field was discarded.

The local oscillator mode was much larger than the probe mode. This helped reduce fluctuations in the measurement due to beam pointing caused by heating of the air near cells. It also helped prevent the increase in the probe mode size following storage in the memory affecting the measurement.

5.1.2 Experimental protocols

The control field was left on constantly during an experimental run. A series of probe pulses were sent over about 180 μs. The magnetic field gradient was switched twice between each sent pulse, allowing for recalls alternating with each sent pulse, as well as simultaneously with each sent pulse. This allowed the interference between sent and recalled pulses to occur only while the first gradient coil was switched on. This improved the interference effect, as the carrier frequency of the pulse is only necessarily preserved when recalling with the same gradient. If the opposite gradient is used and the carrier frequency of the stored pulse does not match the central frequency of the memory, the position offset from the centre of the memory is translated into a frequency offset of the recalled pulse. The timings for the fields are shown in Fig. 5.4

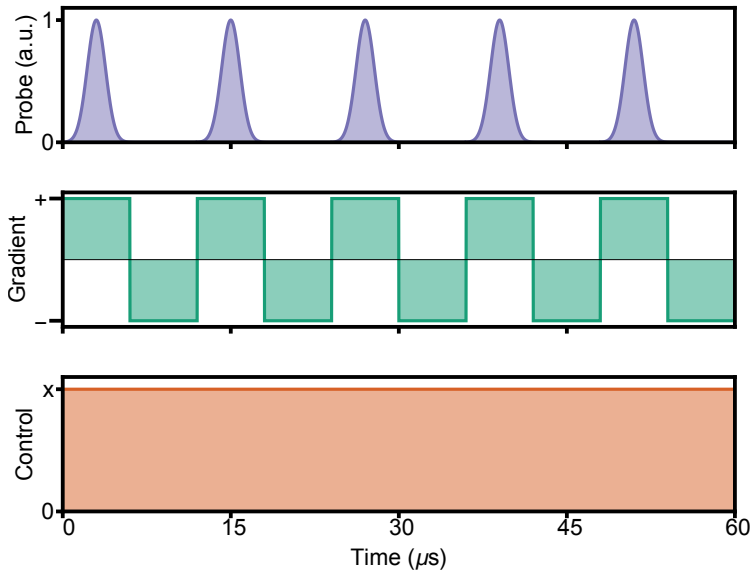

Fig. 5.4. Example timing scheme for the experimental fields. The control field intensity is changed between datasets to demonstrate the dependence of the resonator linewidth on the coupling, and the repetition rate is also varied. Pulses are always sent with a positive gradient.

The control field frequency was scanned by 800 kHz over 200 s. This produced an increment in the phase of the interference between each sent and recalled pulse. When that increment was a multiple of 2π, interference was constructive into the stored mode. For each series of pulses, the control field frequency changed a negligible amount. This allowed the measurement at each frequency of the constructive interference. Some raw data is shown in Fig. 5.5

Fig. 5.6 shows the transmission of a continuous probe sent while scanning the control field detuning and without switching the gradient. Only about 10% of the probe is absorbed inside the memory bandwidth. The large dip at 500 kHz two-photon detuning, and the oscillation in transmission to the right of that point was likely caused by a flat spot in the gradient. This causes a larger number of atoms to absorb light at the same frequency, decreasing the transmission at that frequency. The interference effect was slightly complicated by the diffusion of atoms. The mode of the output probe field was larger than the mode of the input probe field due to the movement of the atoms through the cell. There may also be some small changes in the direction of the output

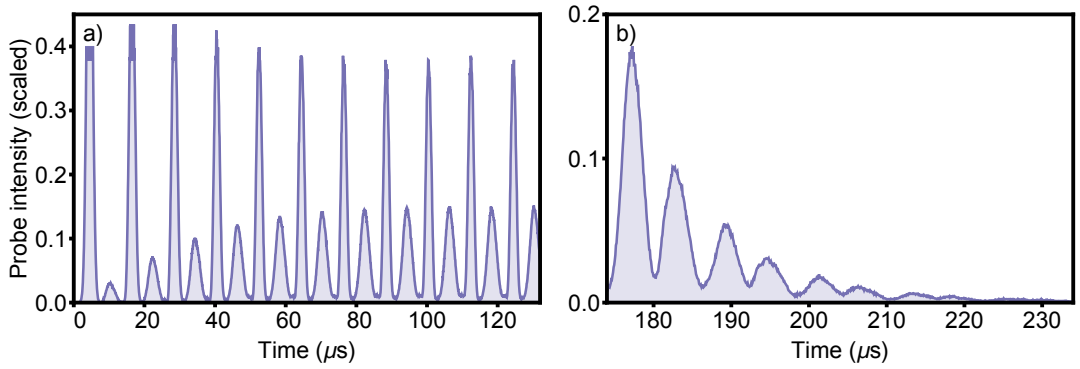

Fig. 5.5. Probe output intensity data scaled to the input pulse intensity for a) the beginning and b) the end of an experimental run. The data is obtained by demodulating the heterodyne measurement. The higher peaks in a) correspond to the leakage of sent pulses and the lower peaks the recalled pulses. The first peak in b) is the first recall after the final probe pulse is sent. The spectra in Figs. 5.7, 5.8. are obtained by charting the height of this peak as a function of control field detuning.

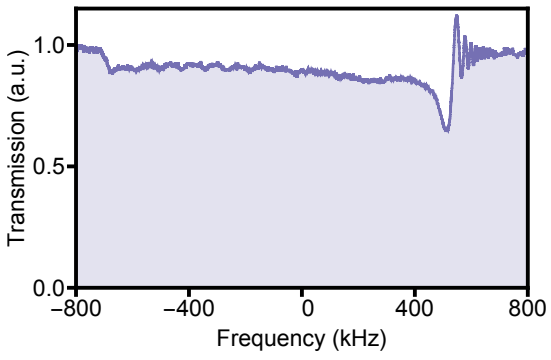

Fig. 5.6. The broadened Raman absorption line with a typical control field, as two-photon detuning is scanned. The dip and oscillation are due to a flat spot in the gradient, where a larger number of atoms absorb at that frequency, and then are allowed to undergo free induction decay.

mode due to imperfect alignment of the magnetic field gradient.

5.2 Results and analysis

The transmission of the system is considered to be the intensity of the steady-state recalled pulse. The system reaches this steady-state after at most 10 pulses, and runs with stronger coupling require fewer pulses to reach the steady-state. A typical resonator response is found as the two-photon detuning is varied of a series of Lorentzian shaped peaks separated by the free spectral range expected from the pulse repetition rate. The response is multiplied by the broader roughly Gaussian envelope related to the bandwidth of the pulse compared to the bandwidth of the memory. The envelope is proportional to the efficiency of storage followed by recall of a pulse in the memory, dependent on the central frequency of the pulse. That is why this envelope appears much narrower than the memory bandwidth expected from the Raman absorption measurement. The spectra for various repetition rates and narrow linewidths are shown in Figs. 5.7,5.8.

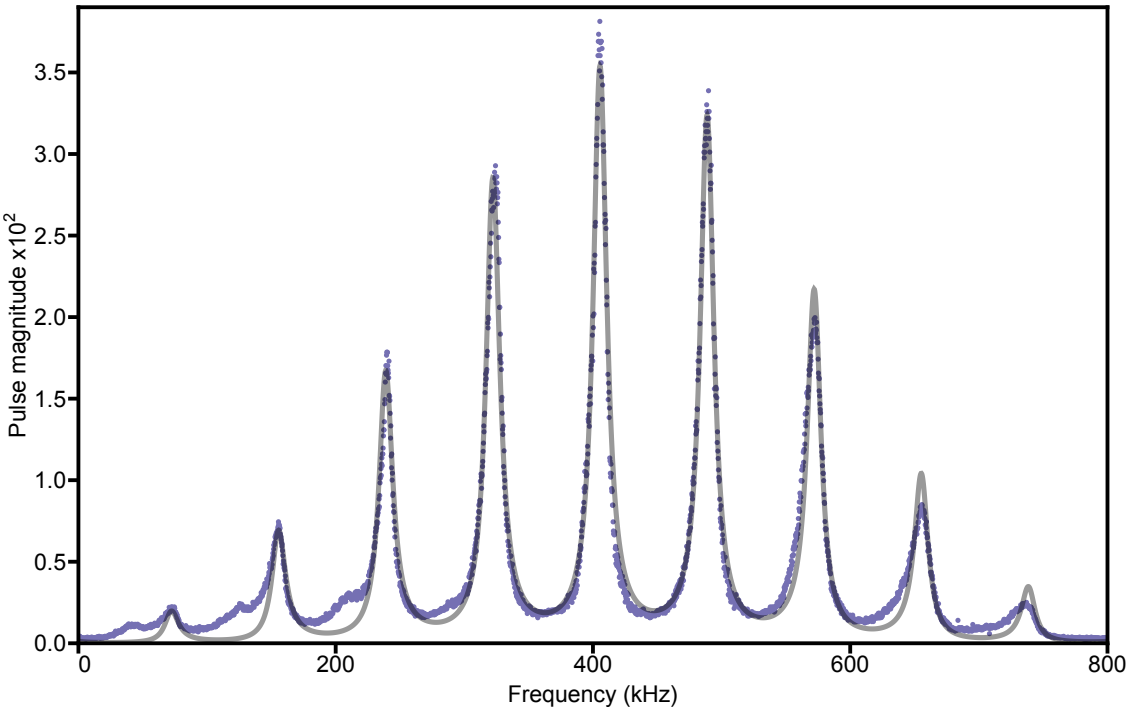

Fig. 5.7. A typical complete spectrum for the spinwave resonator with an input pulse repetition rate of $83.\dot{3}$ kHz. The magnitude of the first pulse recalled after all the sent pulses is measured as the frequency of the control field is scanned. This is plotted as the ratio of this magnitude against the input pulse magnitude. The fit (grey) is constrained to the measured linewidth and FSR, and convolves a square absorption window with a Gaussian pulse bandwidth to fit the broader envelope.

The steady-state data displays a slightly different effective linewidth to that being measured from the decay. This is due to the spinwave mode changing over subsequent pulses. This does not affect the decay of the retrieved pulses once the sent pulses finish.

Fig. 5.9 includes the processed data about the resonator response of the system. The steady-state recall intensity is proportional to the spinwave magnitude (for a single con-

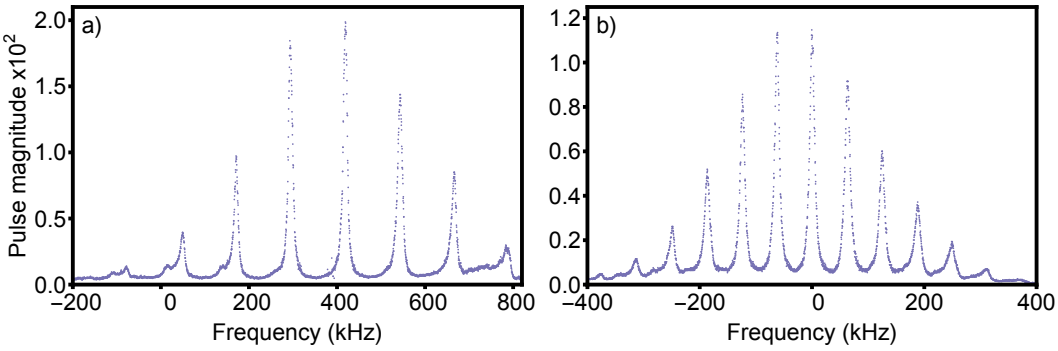

Fig. 5.8. Spectra for input repetition rate of a) 125 kHz and b) 62.5 kHz. The envelope shifts between spectra as the frequency gradients are adjusted to the pulse bandwidths.

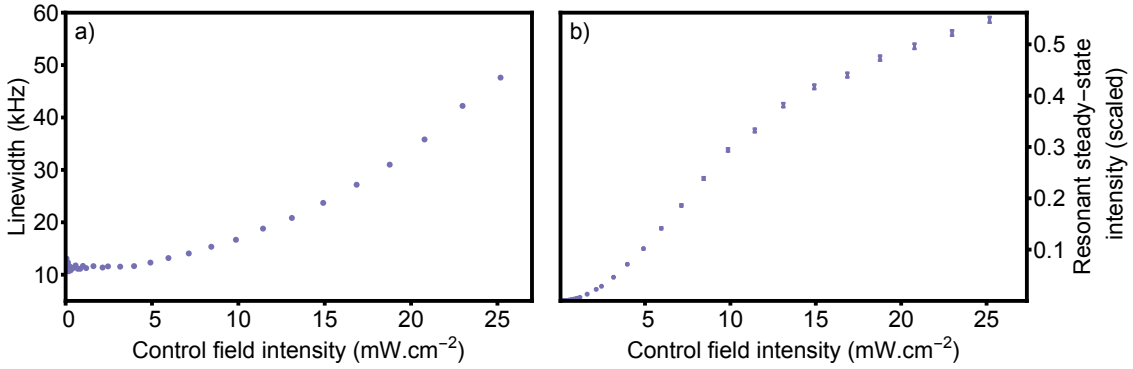

Fig. 5.9. Resonator response as a function of control field intensity. The steady-state recall pulse a) linewidth fitted by Airy transmittance function, and b) maximum recall intensity as a fraction of the input pulse intensity.

trol field intensity - the change in control field intensity changes the ratio between the spinwave magnitude and the steady-state output). This is then used to measure the linewidth of the system. Due to the non-exponential nature of some loss processes, the linewidth as a function of coupling strength cannot be fitted exactly with a typical model for the transmittance of a resonator. Instead, a phenomenological model estimating these losses in the spinwave is used. The model is of the form

$$I(\delta,\Omega)=\sum_{n=1}^{15}T(\Omega)\exp(-n(\alpha(\Omega)+i\delta))\exp\left(-(n\beta)^2\right). \tag{5.1}$$

with the function $T(\Omega)$ the recall efficiency dependent on the control field, and the parameters $\alpha(\Omega)$ and β corresponding to the exponential loss rate (which is control field dependent) and the Gaussian loss rate respectively.

The pulses during the interference phase are recalled in FIFO configuration, so each loss process applies uniformly across the envelope of the recalled pulse. This allows the intensities of the interfering pulses to be described as the above sum. The calculated linewidth is the full width at half maximum of the fitted function.

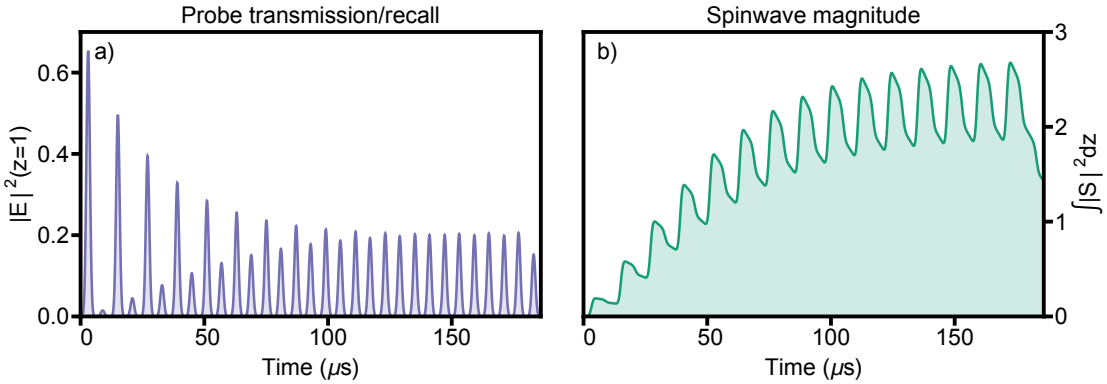

Fig. 5.10. Simulation results for the spinwave resonator. a) The probe intensity at the output of the memory, scaled to the input pulse. b) The spinwave magnitude, integrated across the memory, scaled to the efficient storage of a single pulse.

Fig. 5.10 shows the results of simulations of the spinwave resonator. One important difference between simulation and experimental results is that at steady-state, the transmission intensity while the input probe is sent is roughly equal to the intensity of the recalled pulses. This is because the transverse mode of the spinwave in the simulations does not change compared to the mode of the input, as this effect cannot be included in a one-dimensional simulation. The simulation suggests that the stored mode constructively adds to about 2.8 times the efficient storage of single pulse. In the experiment this enhancement is likely smaller considering that the recalled pulses have lower energy compared to the input pulse.

5.3 Conclusion

In this chapter we presented a method for operating the gradient echo memory as a pulsed optical resonator. We showed how the resonator characteristics depended on sev-

eral easily tunable parameters. This tunability should allow some more complex applications of the protocol. It would also be interesting to obtain sharper linewidths by using a longer lifetime memory such as in cold atoms.

Chapter 6

Gradient echo memory in one dimension

> Oh! cette double mer du temps et de l'espace
> Où le navire humain toujours passe et repasse,
> Je voulus la sonder, je voulus en toucher
> Le sable, y regarder, y fouiller, y chercher,
> Pour vous en rapporter quelque richesse étrange
>
> Victor Hugo, *La pente de la rêverie*

The gradient echo memory relies on an invertible spatial gradient in the resonant frequency of the memory medium. This frequency gradient is typically produced by an external field modifying the energy levels of the medium. As discussed in Sec. 3.4.5, an effective frequency gradient can be achieved by other methods. One proposal involved scanning the incidence angle of the control field to produce a memory with similar characteristics to GEM [78]. Building an efficient memory by control field angular scanning relies on a memory much longer than it is wide. The change in the transverse momentum of the control field as the angle is scanned negatively affects the memory coupling, and the dependence on the memory dimensions is examined in detail in Zhang et al. [78]. An ideal memory with control field angular scanning would have a single spatial dimension. Fig. 6.1 illustrates how control fields Ω_c and Ω_c' sent at different angles couple the probe field $\hat{\mathcal{E}}$ to different momentum vectors of the spinwave, and how those vectors project in a single dimension.

In this chapter we present two methods that take advantage of a single dimensional memory and angular freedom of the control field: a method for extending GEM by rapidly accessing multiple regions of momentum-space (k-space) without rephasing, and a method for synthesising an angular-scanned control field. We then investigate the possibility of implementing one of these schemes in a one-dimensional memory using caesium trapped in the evanescent field of a nanofibre.

Compared to an angular scanned control field, a frequency gradient has the advantage of being able to transport the spinwave indefinitely through k-space. A scanned control field is restricted to the k-space which can be reached in the physical implementation of the scheme. However, it can be significantly faster to scan the control field to quickly access the region of k-space that contains the wanted information. A frequency gradient is limited in this respect by how large the gradient can be made and how quickly it can be switched.

We imagine rapidly time-multiplexing useful photonic states or using the memory

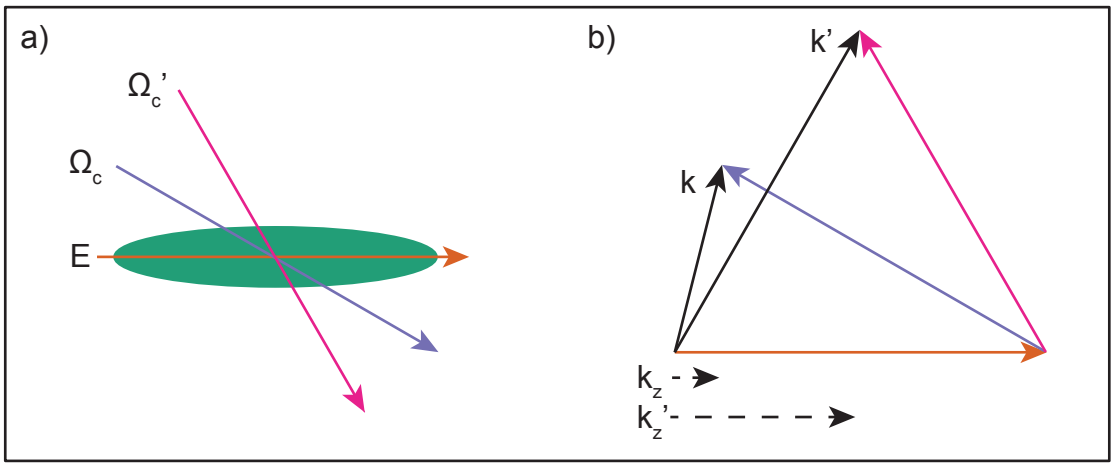

Fig. 6.1. a) Illustration of two control fields Ω_c and Ω_c' coupling the field E into/out of a spinwave in an ensemble. b) The control fields couple to spinwaves with momentum $\vec{k}$ and $\vec{k}'$ respectively. For a one-dimensional memory, only the projections onto the memory dimension z, k_z and k_z', are conserved.

as a random-access quantum memory by adding a number of control fields at different angles to access different momenta. If a sufficient number of momenta can be accessed, a scheme comparable to the angular scanning of the control field can be implemented. Angular scanning of the field can be synthesised by sending an appropriate time-dependent coupling power with each control field.

6.1 Momentum-space access

It may be impractical to angularly scan a single control field through the desired momentum space. In this case, static control fields accessing different points in k-space can be used in addition to a frequency gradient. The frequency gradient transports the spinwave through k-space by applying the spatially varying phase. The different control fields couple to spinwaves at different points in k-space. A spinwave can therefore be stored using one control field and retrieved using another, either without having to flip the gradient, or without having to wait as long after flipping the gradient as is necessary with a single control field. This concept is illustrated in Fig. 6.2. This allows significant flexibility in pulse sequencing, depending on the number of points that can be accessed. The on-demand nature of the memory is improved over GEM if control field switching is quicker than the transport of the spinwave through momentum space by the frequency gradient. For example, in Fig. 6.2 FIFO retrieval could be performed immediately after the first five pulses are stored by turning off control field (i) and turning on control field (ii).

6.2 Momentum-space synthesis

If a sufficient number of control fields can be sent into the memory from different angles, it is possible to synthesise a continual scan of the coupling field through k-space. This is done by adjusting the intensity of the control fields with time. We choose a basis for the

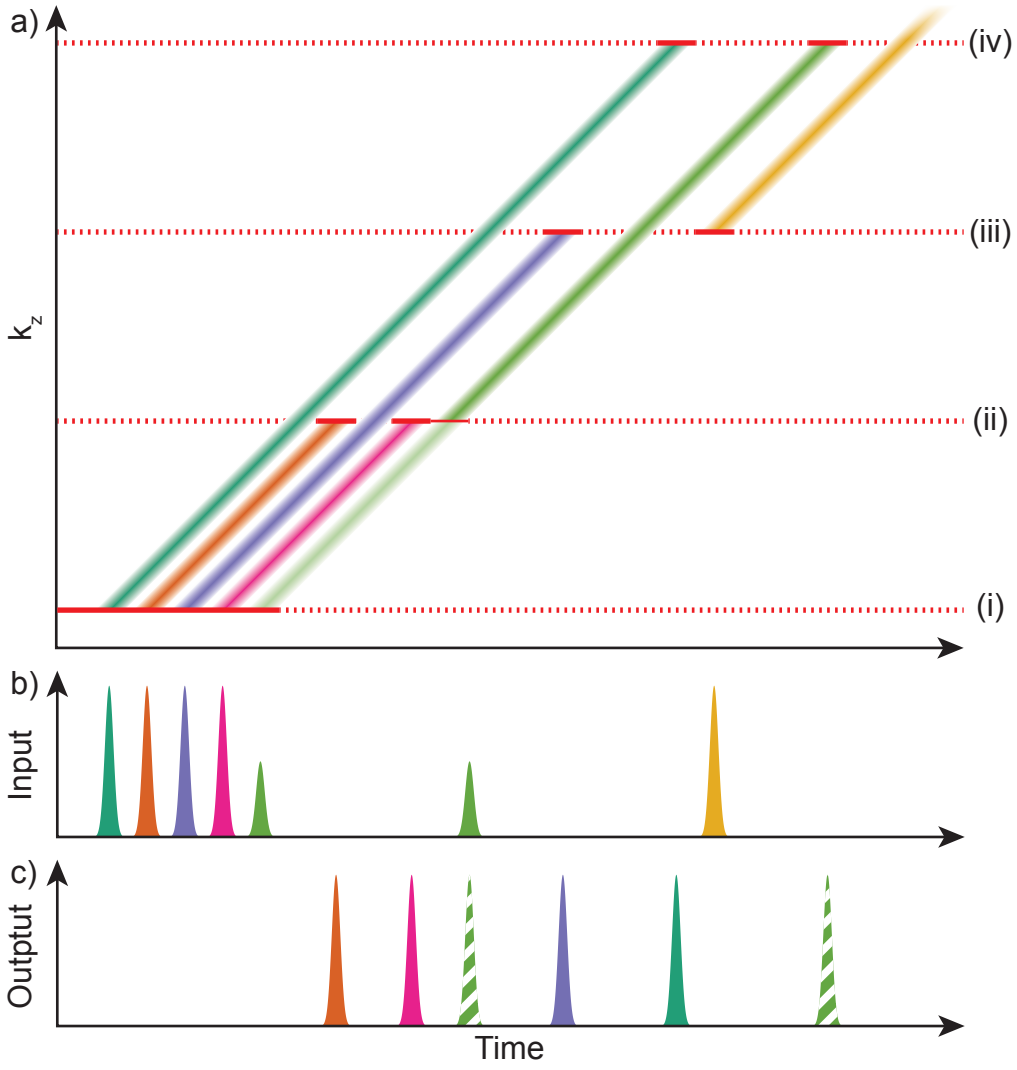

Fig. 6.2. Concept diagram for storage and recall including re-ordering and beamsplitter operation for a number of optical pulses in a k-space access memory. a) The creation/destruction of the spinwaves and their transport through the reduced k-space. The dotted lines (i-iv) show the points in k-space that are coupled by separate control fields. Solid red lines show the times at which each control field is sent. b) input and c) output intensities of the optical pulses. The beamsplitter operation is performed by sending a reduced control field power while simultaneously retrieving and sending pulses. The striped output pulse areas are dependent on the phase and amplitude of the interference between the two pulses. The spinwave (yellow) transported past the last coupling point (iv) is only accessible by reversing the gradient and thus the direction of transport.

control fields by separating them in k-space by units of $k = 2\pi n$. The appropriate time profile for a control field $\Omega_n(t, \xi)$ corresponding to coupling the momentum $2\pi n$ will then correspond to a term in the Fourier series

$$\Omega(t, \xi) = e^{i2\pi \mathrm{BW} t.2\pi\xi} = \sum_{n=-\infty}^{\infty} \mathrm{sinc}(2\pi \mathrm{BW}(t - \frac{n}{2})) \exp(i2\pi n\xi), \tag{6.1}$$

with the longitudinal dimension of the memory ξ and the desired bandwidth BW. We will show with simulations that this arrangement allows an effective memory even for a small number of control fields and storing a small number of modes.

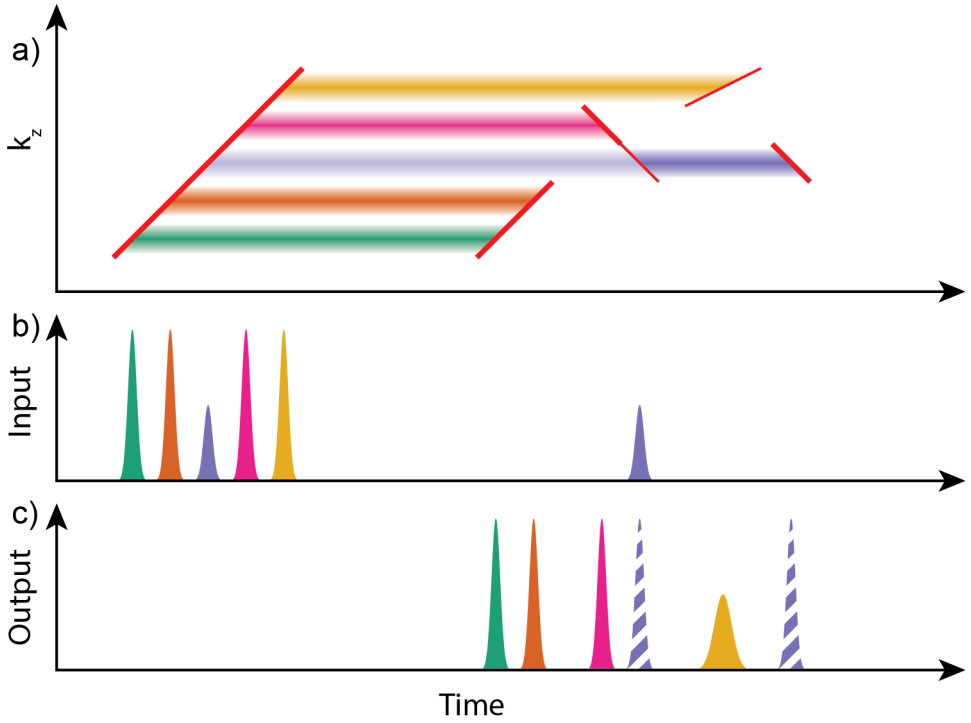

Fig. 6.3. Concept diagram for an angular-scanned or -synthesized memory. a) The writing and reading out of the spinwaves. b) input and c) output intensities of the optical pulses. FIFO and FILO storage, reordering, beamsplitter operation (with relative intensities of the striped output pulses depending on the phase of the interfering spinwave and input pulse) and spectral reshaping are demonstrated. The timing and momentum-space are more compact than in the momentum-space access scheme due to the ability to quickly access a particular momentum.

The synthesised gradient scheme has the same advantage as the angular scanning technique of allowing the implementation of a gradient echo memory on ensemble transitions that are insensitive to magnetic or electric fields. This allows for the same benefits in long storage times and lower phase noise. Pulse reordering and other possible operations of the scheme are shown in Fig. 6.3

A major technical challenge to implementing this scheme is the requirement that the phase of each control field not only be held stable with respect to all the other control fields, but at the correct phase so that the bright fringe of the wavelength-scale spatial interference pattern occurs on the axis of the memory. Otherwise the strength of the coupling will vary greatly with time, reducing the effectiveness of the memory.

Fig. 6.4 plots the momentum-space synthesis for a range of control field numbers. The momentum space coupling converges to the function $sinc(\pi(k_z - t))$ rather than the Dirac delta function $\delta(k_z - t)$ because the memory has a finite length. With larger numbers of control fields, more modes can be stored and the storage becomes more reliable.

Operations such as frequency manipulation could be performed by dividing the memory up into several memories along its length. This can be done by adding additional sets of control fields where the angle changes around a different longitudinal position.

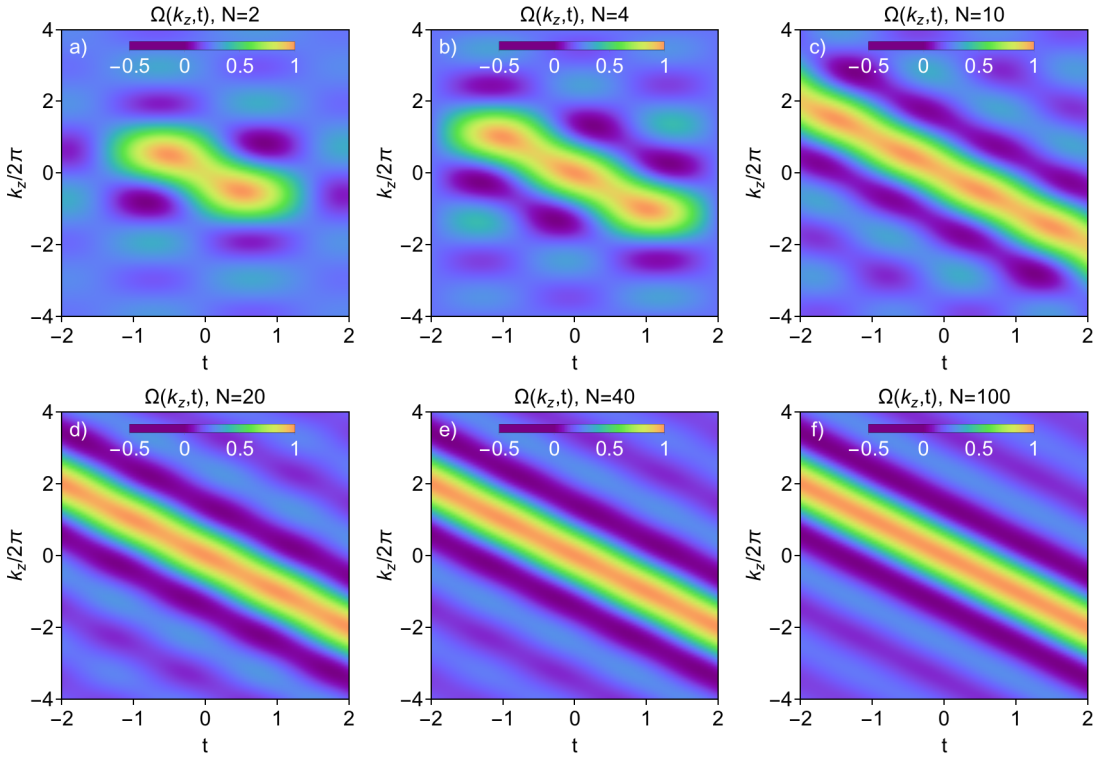

Fig. 6.4. A representation of the momentum-space accessed by a scanned control field synthesised from different numbers N of individual control fields. Due to the finite length of the memory, the control field in momentum-space approaches the function $\Omega(k_z, t) = \mathrm{sinc}(\pi(t - k))$.

Fig. 6.5 shows probe and spinwave behaviour for simulations of FILO retrieval of a pulse with varying numbers of control fields. The control field time profiles are arranged so that the intensity peak of the input and output probe pulses are at the midpoint between the intensity peaks of two control fields. This has the least variance of control field power at low control field numbers. The control field intensity and bandwidth are optimised for the input pulse, and there is not much efficiency advantage by adding more than two control fields.

The simulations in Fig. 6.6 have two input pulses. The control field time profile is shifted so that the midpoint between two control fields lies at the midpoint between the input pulses. FILO retrieval is done by time reversing the control field profile, while FIFO retrieval is done by repeating the control field time profile. The efficiency is slightly lower than optimal as the same control field intensity is used as for single pulse storage. The arrangement of the inputs compared to the control field time profile is optimal for this

input, but even two control fields are sufficient for distinguishing the two inputs.

A more detailed analysis of the multimode capability of this memory is difficult as the scheme lies somewhere between a gradient echo memory and a Raman memory in terms of multimodality. Both these types of memory are very well characterised by different methods, making it difficult to adapt an analysis to the middle ground. Given the difficulty of implementing momentum-space synthesis with large numbers of control fields, a simpler version could involve small numbers of control fields well separated in momentum space to create a multimode Raman memory.

6.3 Nanofibre experiment for momentum-space access

A memory based on atoms trapped in the evanescent field of a nanofibre is approximately one-dimensional, and truly one-dimensional when a single row of atoms is trapped or coupled next to a straight fibre. Momentum-space access and scanning schemes are therefore possible without having to consider the effect of transverse momentum on the memory efficiency. We designed an experiment based on the concept of momentum-space access, with separate control fields used for storing and retrieving a probe pulse. A constant longitudinal magnetic field gradient was applied and not switched between storage and recall. Instead, the angles of the control fields were chosen so that the read control field coupled to the spinwave momentum of the written spinwave plus the momentum added by the gradient during storage.

This scheme was particularly useful for the experimental setup in which it was implemented as the gradient switching time was quite long compared to the expected lifetime of the memory. By removing the requirement to switch the gradient coils, a larger storage bandwidth could also be used. This would allow a pulse to be recalled before the efficiency dropped too far.

In a nanofibre experiment the available polarisations for the probe beam are restricted. The probe field must travel through the nanofibre to achieve an enhanced optical depth due to the evanescent coupling to the atoms, and the polarisation of the evanescent field is restricted due to the nanofibre mode. The polarisation for a single-mode nanofibre is quasi-linear. This does not mean the polarisation is roughly linear but describes the distribution of the probe polarisation within the nanofibre. For a quantisation axis parallel to the fibre axis, the polarisation evanescent field has a significant pi component due to the steep amplitude gradient of the field in the radial direction.

The restricted fibre polarisation places additional restrictions on the memory scheme. If the polarisation emitted by the atoms differs from that of the evanescent field, the polarisation in the fibre is perturbed. This puts some portion of the field in a mode not supported by the fibre, causing loss. The amount of loss may depend on the length scale over which the perturbation occurs. It is unclear how this would affect the efficiency of GEM in this system. The absorption length could be tuned and the atomic coupling changed under any system of atom-light coupling, allowing study of efficiency limits and how they might impact different memory schemes.

Pointing the magnetic field along the gradient will make the quantisation axis point along the fibre axis. The polarisation of the evanescent field with respect to the atoms is then a mixture of circular and π polarisation. Generating a linear magnetic field gradient of sufficient magnitude with the field not pointing along the direction of the gradient is more difficult.

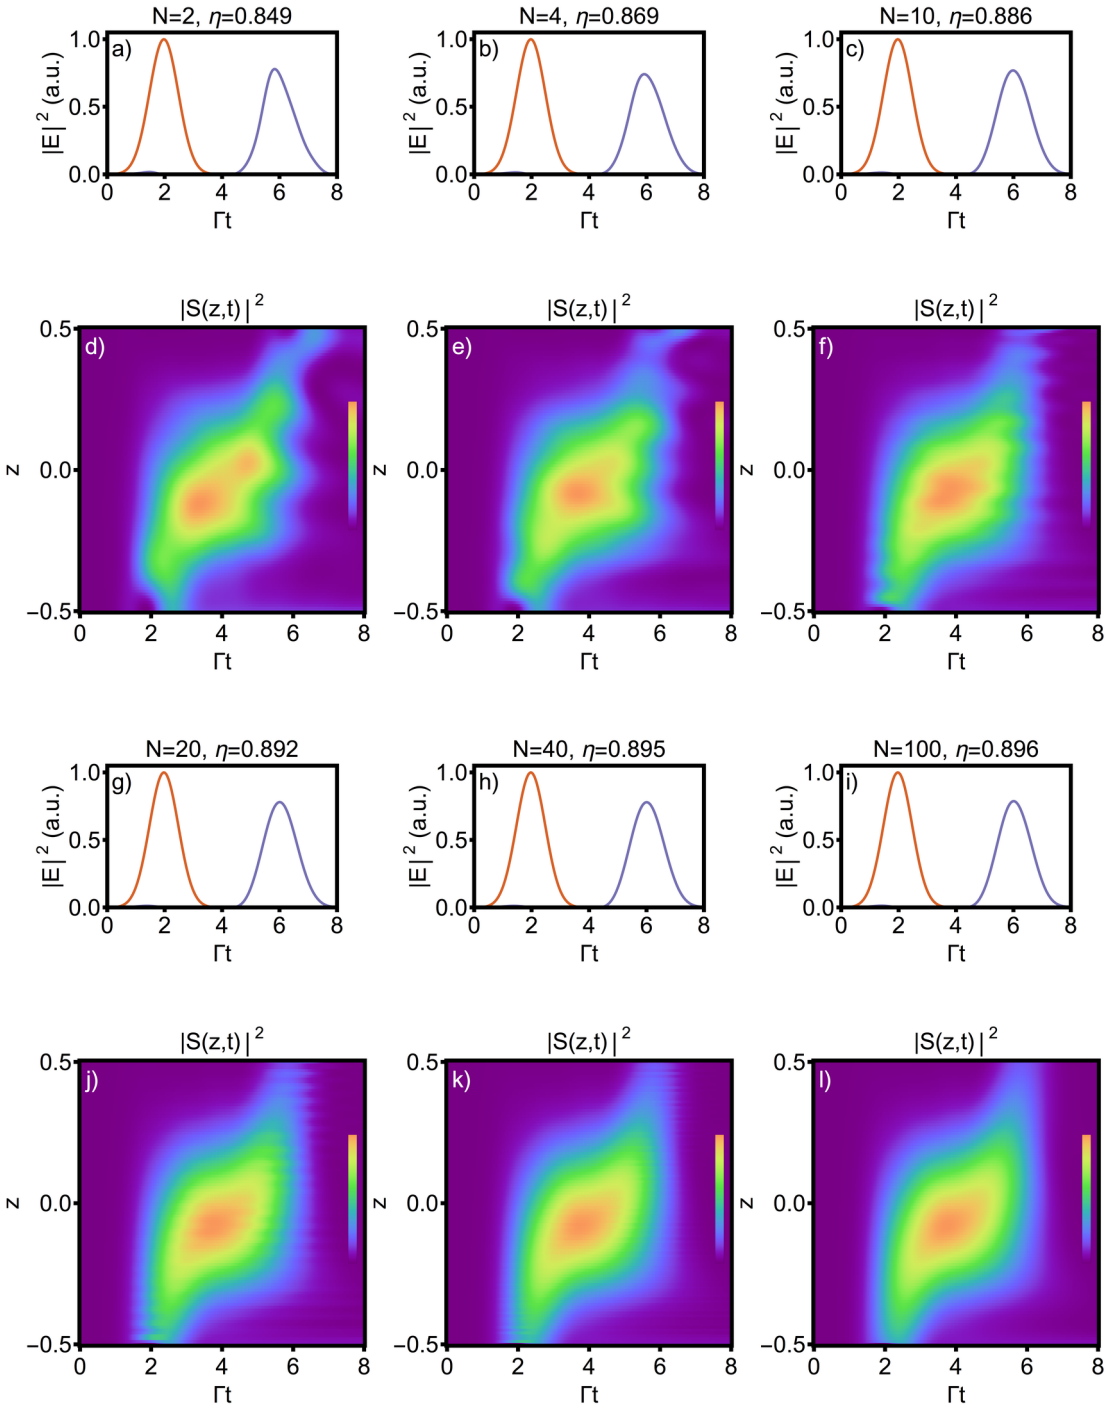

Fig. 6.5. Simulations of synthesised momentum-space access. Plots a)-c),g)-i) show input (orange) and output (blue) intensities, labelled with control field number N and total retrieval efficiency η. Plots d)-f), j)-l) show corresponding spinwave magnitudes through the memory.

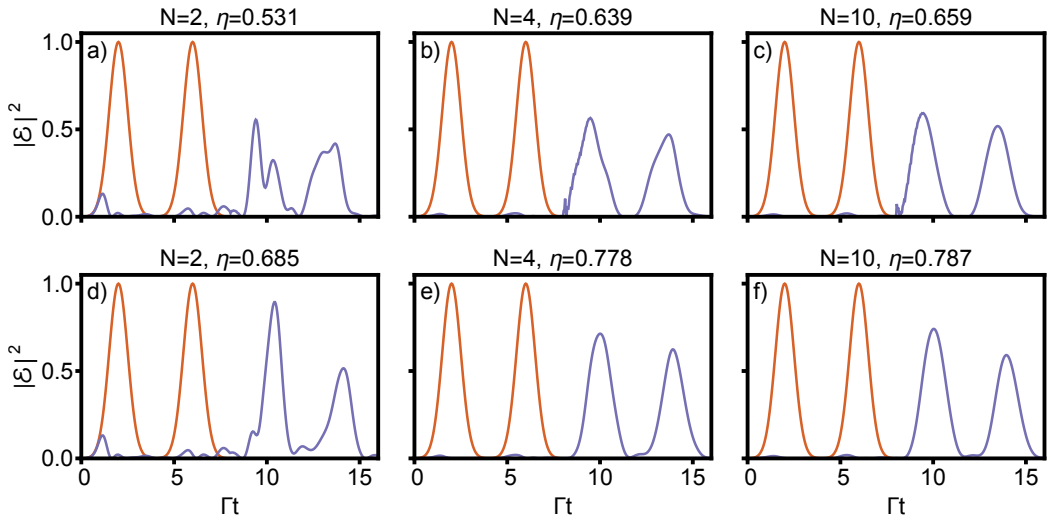

Fig. 6.6. Simulations of synthesised momentum-space access. Plots a)-c) show input (orange) and output (blue) intensities for FIFO retrieval, labelled with control field number N and total retrieval efficiency η. Plots d)-f), show corresponding plots for FILO retrieval.

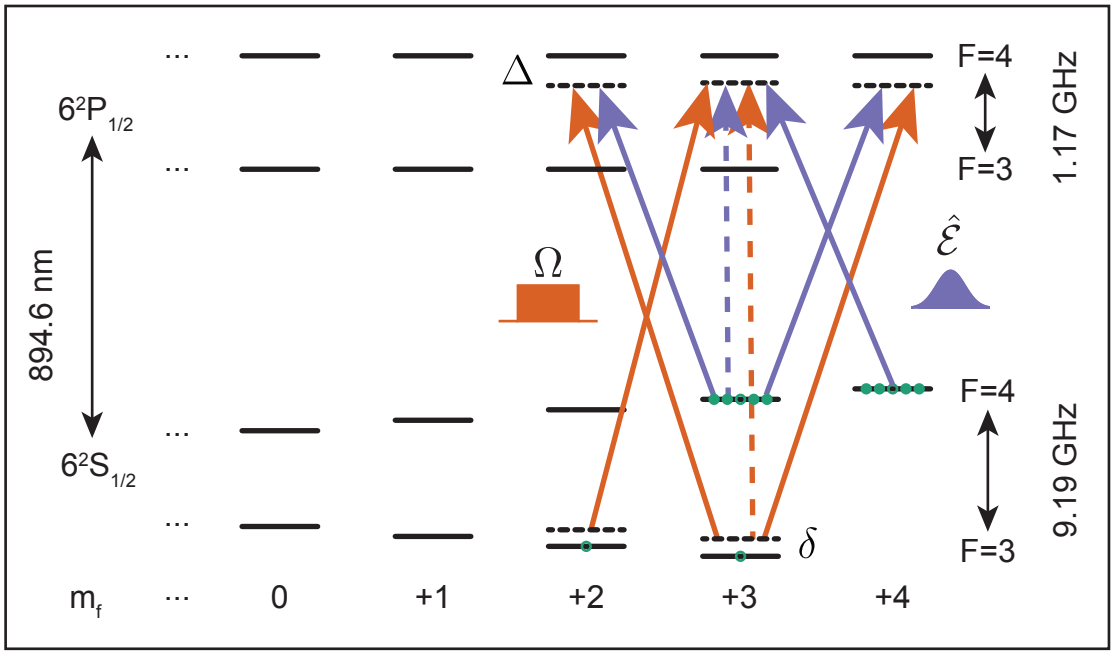

Fig. 6.7. Level scheme for the caesium D1 line, with experimental fields indicated. The restricted polarisation in the nanofibre requires the use of two populated levels to balance the coupling for each circular polarisation. Adding π polarised control field allows coupling of all the probe polarisations (dashed lines). Atomic energy levels taken from Steck alkali data [114].

The level scheme for the caesium D_1 line and the transitions used in the memory are shown in Fig. 6.7. The Raman transition must be chosen carefully. The edge state $|F = 4, m_f = +4\rangle$ cannot be used on its own, as one of the circular polarisations will not be absorbed, greatly reducing the efficiency. Using $|F = 4, m_f = +3\rangle$ as well as $|F = 4, m_f = +4\rangle$ is more effective. The transitions $|F = 4, m_f = +3\rangle \rightarrow |F' = 4, m_f = +2\rangle$ and $|F = 4, m_f = +3\rangle \rightarrow |F' = 4, m_f = +4\rangle$ have quite different coupling strengths. However, the extra transition $|F = 4, m_f = +4\rangle \rightarrow |F' = 4, m_f = +3\rangle$ helps balance the effective coupling strength for each polarisation. The coupling can be further balanced by changing the relative populations of the two initial states. Population in other ground states does not affect the memory coupling as there are no transitions on two-photon resonance. At higher optical depths those states will contribute to incoherent absorption. Pumping into the coupled levels will then be useful beyond simply increasing the coupling.

The relative coupling for the two polarisations, assuming the two ground states are equally populated and the control field is linearly polarised (equal mix of σ_+ and σ_-), is calculated by multiplying the Clebsch-Gordon coefficients for each pair of transitions. These coefficients are taken from [114]. This calculation follows the notation of [74], where $\mu_{(m,n;p,q)}$ is the dipole matrix element for the transition $|F = m, m_f = n\rangle \rightarrow |F' = p, m_f = q\rangle$. We ignore the small difference in one-photon detuning Δ for each transition caused by the Zeeman splitting.

$$g_- \propto \mu_{(4,4;4,3)}\mu_{(3,2;4,3)} + \mu_{(4,3;4,2)}\mu_{(3,3;4,2)} = 0.121 \tag{6.2}$$

$$g_+ \propto \mu_{(4,3;4,4)}\mu_{(3,3;4,4)} = 0.110 \tag{6.3}$$

At higher m_f numbers, the atomic frequency gradient is a larger multiple of the magnetic field gradient. The memory bandwidth is then larger for a given magnetic field gradient, and a higher time-bandwidth product and efficiency are possible.

The absence of coupling for the transverse polarised field reduces the efficiency of the memory. The coherent emission from the atoms is in a particular eigenmode which only includes the circular polarisations, and therefore the transverse component does not interact. This remaining component has a low overlap with the single mode of the nanofibre and is mostly lost. When the pulse is recalled, the atoms again emit only in circular polarisation. The overlap between the atomic eigenmode and the nanofibre mode is not perfect and some of the light is not coupled back into the fibre. This loss can be prevented by introducing a transverse component in the control field polarisation, and choosing an appropriate two-photon transition. For example, with the chosen two-photon transition, the addition of a transverse component couples $|F = 4, m_f = +3\rangle$ to $|F = 3, m_f = +3\rangle$, the same transition as for the two circularly polarised pairs interacting with this level.

The effect on the efficiency of GEM of having only a few atoms within a given two-photon linewidth has not been calculated. This could reduce the effective optical depth, and make a GEM protocol impossible for the given experimental parameters.

6.3.1 Experimental setup

A good account of the use of nanofibres and evanescent trapping in atom-optic memories can be found in the thesis of Baptiste Gouraud [42]. The basic concepts are also worth mentioning here. A diagram of a nanofibre and the trapping of atoms in an evanescent

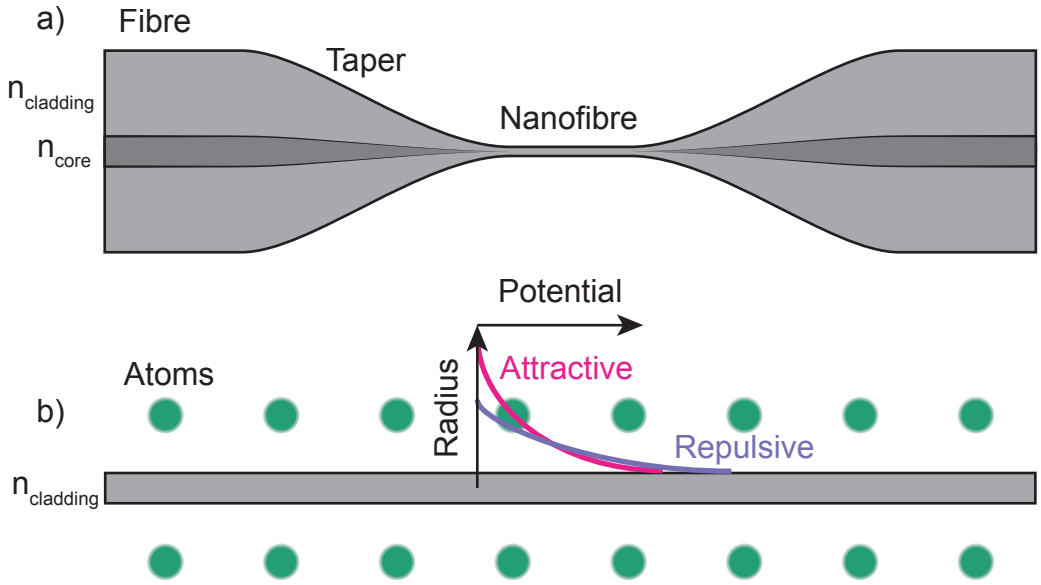

Fig. 6.8. a) The profile of a nanofibre tapered from regular fibre with a core index $\mathrm{n_{core}}$ and cladding index $\mathrm{n_{cladding}}$. b) A small section of nanofibre, with atoms (not to scale) and evanescent trapping fields. The trapping point is the distance from the nanofibre at which the attractive and repulsive forces due to the two trapping fields cancel. The evanescent field of the probe extends in a similar way from the nanofibre.

field are shown in Fig. 6.8

The atoms were first trapped and cooled in a MOT. Polarisation gradient cooling then cooled the atoms enough that they could be trapped in the evanescent field of the fibre. Trapping lasers were sent into the fibre to produce a dipole trap in the evanescent field. By shining a blue detuned (repulsive) and red detuned (attractive) laser, the atoms were held some distance from the fibre. By choosing appropriate detunings, the light-shift (AC-Stark) of the two-photon resonance due to the trapping fields was cancelled out. The trapping lasers were sent in each direction along the nanofibre. The red detuned lasers were at the same frequency and formed a standing wave, while the blue detuned lasers were at different frequencies. The average potential for the atoms formed a series of wells along two opposite sides of the nanofibre, at intervals of half the red detuned laser's wavelength.

An optical depth of about 100 was achievable in that setup [115], considering atoms in all m_f states of the $|F = 4\rangle$. This corresponded to an estimated optical depth of about 10 for the Raman transition accessed by the probe field in the effective 3-level ensemble. This also corresponded to a population of several hundred atoms in each of $|F = 4, m_f = 3\rangle$ and $|F = 4, m_f = 4\rangle$.

We attempted to demonstrate a k-space access memory by storing a pulse using one control field and a magnetic field gradient. The pulse $\hat{\mathcal{E}}$ and control field Ω follow the level scheme in Fig. 6.7, with a detuning of $\Delta = 50$ MHz. A second control field at a larger angle to the nanofibre was used to retrieve the stored pulse without flipping the gradient.

Assuming no other unaccounted for effects were present, simulations of the scheme indicated that an experimental measurement of the memory effect should have been pos-

sible. We were able to observe Raman absorption of the probe under experimental conditions, but no echo could be observed. Given that an EIT memory has previously been demonstrated in this system [40], it will be interesting to find out if there is some unexpected problem with implementing GEM in a nanofibre-trapped ensemble.

6.4 Conclusion

We proposed two schemes for increasing the flexibility of a gradient echo memory. These schemes move some or all the control of the memory into the optical fields, potentially allowing for simpler or more robust memories.

We attempted to demonstrate one of the schemes in an atomic ensemble trapped in the evanescent field of a nanofibre. The experiment did not work well enough to test the effectiveness of the scheme, but we were able to show that the scheme should have significant advantages for implementing a GEM-like memory in systems where switching of the atomic frequency gradient is a limiting factor.

The open question of how polarisation-selective absorption in the evanescent field of the nanofibre affects propagation of the light in the nanofibre could be explored in this experimental system. This could easily be achieved by measuring the response of the system when altering the relative coupling of the different polarisations to trapped atoms and adjusting the lengths over which the absorption takes place.

Part II

Stationary light in atomic ensembles

Chapter 7

Theory and literature review of stationary light

Stationary light (SL) is a form of light proposed by Andre and Lukin in 2002 [116] and involves trapping light with EIT by adding an additional counterpropagating control field. The light can then be made stationary without reducing its intensity, increasing the potential for generating nonlinear effects. Stationary light was demonstrated in warm atomic vapour shortly after this proposal [117]. Stationary light has been studied further in warm and cold atoms, and the complex behaviours resulting from counterpropagating optical fields have been explored in detail.

The first proposal involved a spatially modulated light-shift applied to the electromagnetically-induced transparency by applying a standing wave interacting between the state $|2\rangle$ and a fourth level (the level scheme is illustrated in Fig. 7.2 d) [116]. This creates an absorptive Bragg grating by spatially modulating the transparency frequency, trapping the light. In the first demonstration a standing wave was generated in the control field intensity itself [117]. This was also thought to produce a standing wave in the probe absorption due to the spatial modulation of the EIT effect (see Fig. 7.1). Reflection from an atomic vapour illuminated by counterpropagating control fields was also demonstrated around the same time [118].

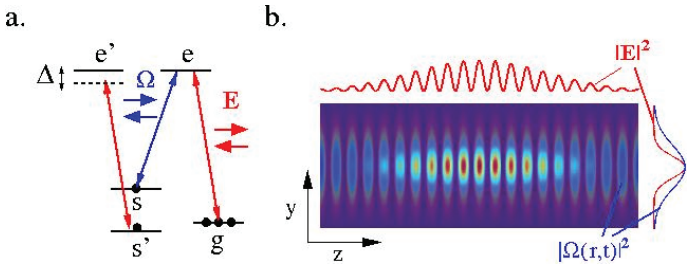

Fig. 7.1. Stationary light generated by standing wave control fields in EIT. a) The EIT scheme with counterpropagating fields, and b) the expected probe intensity due to the grating. The control field Ω creates a transverse confinement due to the absorption toward the edges. Figure in [119]

However, the Bragg grating is not a satisfactory explanation for demonstrations of stationary light. Probe light travelling through the medium does not actually experience a spatially modulated absorption profile. Instead, the coupling to the excited state and subsequent coupling by a control field to the hyperfine state generates slow light. Adding a second counterpropagating control modifies the slow light effect, but does not directly

interact with the original probe field. This can be understood by the speed of the EIT interaction, which is much slower than the motion of atoms across a period of the spatial modulation of the control fields. All subsequent demonstrations of stationary light have been in systems of moving atoms, and so another effect is necessary to explain how they work.

Stationary light has generally been studied in EIT with the addition of a counterpropagating control field. Complete explanations of the effect involve multi-wave mixing, where a spinwave couples to probe fields travelling in opposite directions. To help in understanding these studies of stationary light, we will first lay out the theory of how light interacts with atoms under the additional field. Parts of this theory can also be found in the following publications:
"**Dynamical observations of self-stabilizing stationary light** J. L. Everett, G. T. Campbell, Y.-W. Cho, P. Vernaz-Gris, D. B. Higginbottom, O. Pinel, N. P. Robins, P. K. Lam, B. C. Buchler, Nature Physics 13(1), 68-73 (2016)"

"**Direct Imaging of Slow, Stored, and Stationary EIT Polaritons** G. T. Campbell, Y.-W. Cho, J. Su, J. L. Everett, N. P. Robins, P. K. Lam, B. C. Buchler, Quantum Science and Technology **2** 034010 (2017)"

7.1 Interaction of three-level atoms with counterpropagating optical fields

We will first find a general expression for how counterpropagating control and probe fields interact with three-level atoms. We again follow the standard derivation [7], starting with the quantised field equations for interaction of strong and weak optical fields. This derivation mirrors that of 2.4, so we try to only include those steps that are made less obvious due to the addition of the counterpropagating fields.

The operators for the two weak counterpropagating quantum fields and two counterpropagating classical control-fields are

$$\begin{aligned}
\hat{\mathbf{E}}_{p+}(z) &= \epsilon_{p+}\left(\frac{\hbar\omega_{p+}}{4\pi c\epsilon_0 A}\right)^{1/2}\int_{\omega_{p+}} \mathrm{d}\omega\left(\hat{a}_\omega e^{i\omega z/c}+\hat{a}_\omega^\dagger e^{-i\omega z/c}\right),\\
\hat{\mathbf{E}}_{p-}(z) &= \epsilon_{p-}\left(\frac{\hbar\omega_{p-}}{4\pi c\epsilon_0 A}\right)^{1/2}\int_{\omega_{p-}} \mathrm{d}\omega\left(\hat{a}_\omega e^{-i\omega z/c}+\hat{a}_\omega^\dagger e^{i\omega z/c}\right),\\
\mathbf{E}_{c+}(z) &= \epsilon_{c+}\mathcal{E}_{c+}(t-z/c)\cos[\omega_{c+}(t-z/c)],\\
\mathbf{E}_{c-}(z) &= \epsilon_{c-}\mathcal{E}_{c-}(t+z/c)\cos[\omega_{c-}(t+z/c)],
\end{aligned} \tag{7.1}$$

with the quantum optical fields travelling in the increasing and decreasing z directions represented by the operators $\hat{\mathbf{E}}_{p+}(z)$ and $\hat{\mathbf{E}}_{p-}(z)$ respectively. The classical control fields travelling in each direction are similarly $\mathbf{E}_{c+}(z)$ and $\mathbf{E}_{c-}(z)$. The direction of increasing z is called the 'forward direction' and decreasing z the 'backward direction'. The $+$ and $-$ subscripts are always used to denote the direction of travel for the fields. ϵ_{p+} is the polarisation vector for the forward travelling probe field, and so on.

The weak fields each exist in a small bandwidth around a carrier frequency given by $\omega_{p+} = \omega_{13} + \Delta_+ + \delta$, $\omega_{p-} = \omega_{13} + \Delta_- + \delta$, $\omega_{c+} = \omega_{23} + \Delta_+$, $\omega_{c-} = \omega_{23} + \Delta_-$ where the frequency difference of each probe-field $\omega_{p\pm}$ relative to the transition frequency ω_{13}

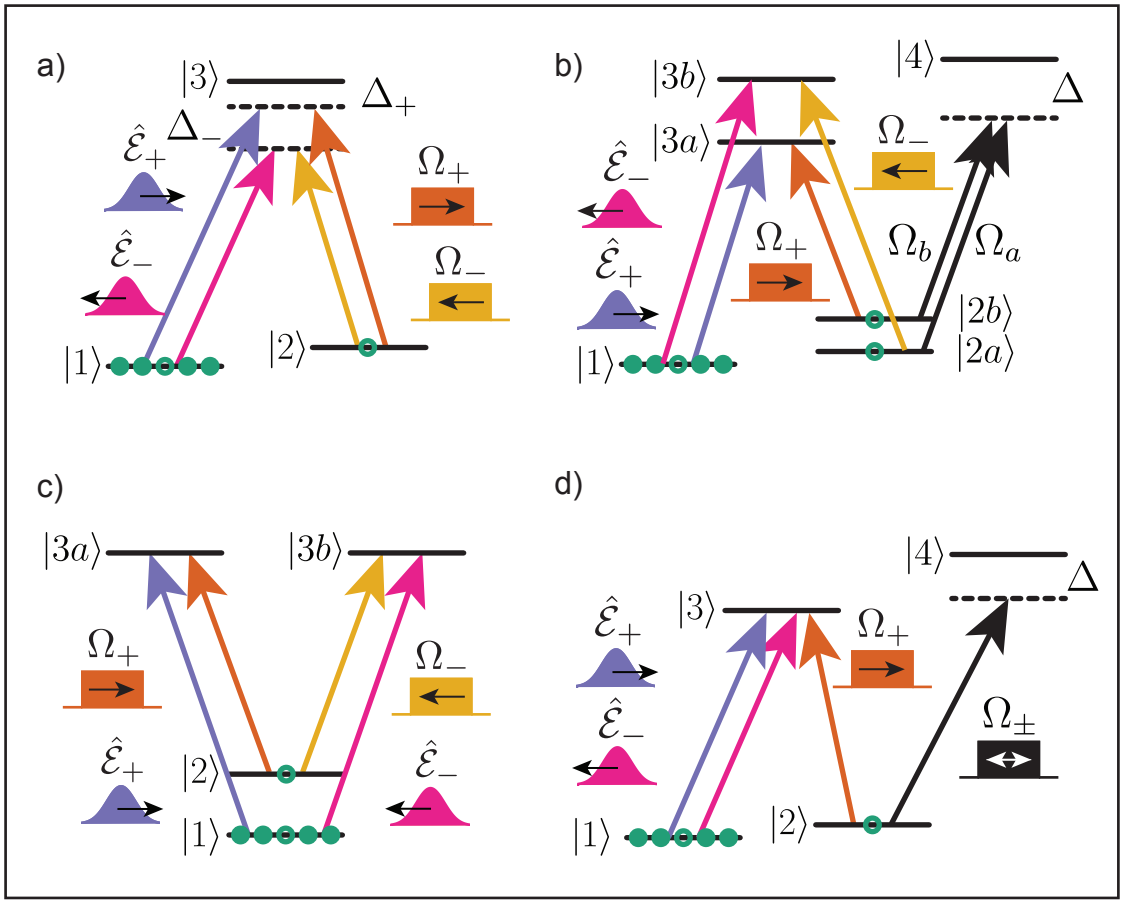

Fig. 7.2. Typical stationary light level schemes. The scheme in a) allows for both multi-wave mixing and Bragg reflection. The scheme in b) is a different form of stationary light that relies on cycling via the additional coupling fields $\Omega_{s\pm}$ between each direction of travel. The scheme in c) is entirely secular and only allows for multi-wave mixing. The scheme in d) is designed purely for Bragg reflection by the modulation of the EIT window by the coupling fields $\Omega_{s\pm}$, but may allow multi-wave mixing.

includes an independent detuning $\Delta_\pm$ from the excited state and a common detuning δ from the ground state. The common detuning δ allows for the non-uniform atomic resonance, as with the simpler co-propagating fields equations. The arrangement of the fields in relation to the atomic levels are illustrated in Fig. 7.2 a).

Writing out only the interaction part of the Hamiltonian gives

$$\begin{aligned}\hat{V} = -\hbar \sum_{i=1}^{N} \Bigg[&\left(\Omega_{c+}(t - z_i/c) e^{-i\omega_{c+}(t-z_i/c)} + \Omega_{c-}(t + z_i/c) e^{-i\omega_{c-}(t+z_i/c)} \right) \hat{\sigma}^i_{32} \\ &+ g \left(\frac{L}{2\pi c} \right)^{1/2} \left(\int_{\omega_{p+}} \mathrm{d}\omega \hat{a}_\omega e^{i\omega z/c} \hat{\sigma}^i_{31} + \int_{\omega_{p-}} \mathrm{d}\omega \hat{a}_\omega e^{-i\omega z/c} \hat{\sigma}^i_{31} \right) + \mathrm{H.c.} \Bigg], \end{aligned} \tag{7.2}$$

where H.c. represents the Hermitian conjugate.

The atomic populations and field envelope over a slice $\mathrm{d}z$ of the ensemble, where the slice is assumed to contain a large number of atoms $N_z \gg 1$, are described by the following operators. The multiplication by terms of the form $\exp(-i\omega(t - z/c))$ creates an individual rotating frame for each operator, causing it to vary only slowly in space and time, and making it useful for calculations. The $\pm$ superscripts indicate that the phase term in the operator is chosen so that the operator is slowly varying with respect to fields travelling in the $\pm z$ direction.

$$\begin{aligned}
\hat{\sigma}_{\mu\mu}(z,t) &= \frac{1}{N_z} \sum_{n}^{N_z} \hat{\sigma}^n_{\mu\mu}(t), \\
\hat{\sigma}^\pm_{32}(z,t) &= \frac{1}{N_z} \sum_{n}^{N_z} \hat{\sigma}^n_{32}(t) e^{-i\omega_{c\pm}(t \mp z_n/c)}, \\
\hat{\sigma}^\pm_{31}(z,t) &= \frac{1}{N_z} \sum_{n}^{N_z} \hat{\sigma}^n_{31}(t) e^{-i\omega_{p\pm}(t \mp z_n/c)}, \\
\hat{\sigma}^\pm_{21}(z,t) &= \frac{1}{N_z} \sum_{n}^{N_z} \hat{\sigma}^n_{21}(t) e^{-i(\omega_{p\pm} - \omega_{c\pm})(t \mp z_n/c)}, \\
\hat{\mathcal{E}}_\pm(z,t) &= \sqrt{\frac{L}{2\pi c}} e^{i\omega_{p\pm}(t \mp z/c)} \int_{\omega_{p\pm}} \mathrm{d}\omega \hat{a}_\omega(t) e^{\pm i\omega z/c}.
\end{aligned} \tag{7.3}$$

The commutators for these collective operators are

$$\begin{aligned}
[\hat{\sigma}_{\mu\nu}(t), \hat{\sigma}_{\alpha\beta}(t)] &= \delta_{\nu\alpha} \hat{\sigma}_{\mu\beta}(t) - \delta_{\mu\beta} \hat{\sigma}_{\alpha\nu}(t), \\
\left[\hat{\mathcal{E}}_\pm(t), \hat{\mathcal{E}}^\dagger_\pm(t) \right] &= 1.
\end{aligned}$$

The slowly varying operators are substituted into $\hat{V}$. The excited state operator $\hat{\sigma}_{33}$ is split into two, $\hat{\sigma}^\pm_{33}$, to allow for the different detunings of the two probes fields. The complete Hamiltonian with the slowly varying operators is

$$\hat{H} = \int \mathrm{d}\omega \hbar\omega \hat{a}^\dagger_\omega \hat{a}_\omega - \frac{\hbar\omega_{p+}}{L}\int_0^L \mathrm{d}z \hat{\mathcal{E}}^\dagger_+ \hat{\mathcal{E}}_+ - \frac{\hbar\omega_{p-}}{L}\int_0^L \mathrm{d}z \hat{\mathcal{E}}^\dagger_- \hat{\mathcal{E}}_-$$
$$+ \int_0^L \mathrm{d}z \hbar n(z) \times \Bigg(\Delta_+ \hat{\sigma}^+_{33} + \Delta_- \hat{\sigma}^-_{33} - \Bigg[\Omega_{c+}(t-z/c)\hat{\sigma}^+_{32} + \Omega_{c-}(t+z/c)\hat{\sigma}^-_{32}$$
$$+ \Omega_{c+}(t-z/c)\hat{\sigma}^-_{32} e^{-i[\omega_{c+}(t-z/c)-\omega_{c-}(t+z/c)]} + \Omega_{c-}(t+z/c)\hat{\sigma}^+_{32} e^{-i[\omega_{c-}(t+z/c)-\omega_{c+}(t-z/c)]}$$
$$+ g\Big(\hat{\mathcal{E}}_+ \hat{\sigma}^+_{31} + \hat{\mathcal{E}}_- \hat{\sigma}^-_{31} + \hat{\mathcal{E}}_+ \hat{\sigma}^-_{31} e^{-i[\omega_{p+}(t-z/c)-\omega_{p-}(t+z/c)]} + \hat{\mathcal{E}}_- \hat{\sigma}^+_{31} e^{-i[\omega_{p-}(t+z/c)-\omega_{p+}(t-z/c)]} \Big)$$
$$+ \mathrm{H.c.} \Bigg] \Bigg), \tag{7.4}$$

where the (z,t) dependence of the operators is omitted for readability. The equations of motion for the operators cannot be found directly from the Hamiltonian as they were for three-level atoms under co-propagating fields in Eq. (2.68). Instead, the cross-terms carrying the rapidly spatially varying phase $\exp(-i[\omega_{c+}(t-z/c) - \omega_{c-}(t+z/c)])$ must first be dealt with. These cross-terms account for the forward travelling probe being coupled to the spinwave by the backward travelling control field, and vice versa, in addition to absorption of control field travelling in one direction and re-emission in the opposite direction. The cross-terms generate coherences with more rapidly varying spatial dependence. In addition, any difference in detuning $\omega_{c+} - \omega_{c-}$ means that the time dependence of the operators is not so simple to write down. This can be accounted for by further dividing up the operators according to spatial and time dependence. The coherences with rapidly varying spatial dependence are called higher order coherences.

7.1.1 Higher order coherences

If the difference in one-photon detunings $\Delta_+ - \Delta_-$ is small, the cross-terms carrying the rapidly spatially varying phase $\exp(-i[\omega_{c+}(t-z/c) - \omega_{c-}(t+z/c)])$ cannot be ignored. Instead, the slowly varying operators can be further broken up into terms accounting for the spatial variation.

For example [120],

$$\hat{\sigma}_{12} = e^{-i\delta t} \sum_{n=-\infty}^{\infty} \hat{\sigma}^{(2n)}_{12} e^{in[(\omega_{c+}+\omega_{c-})z/c + (\omega_{c+}-\omega_{c-})t]} \tag{7.5}$$

where each further term is generated by absorbing a control photon travelling in one direction and emitting it in the other.

The optical coherences generated in between are written

$$\hat{\sigma}_{13} = e^{-i\Delta_+ t + i\omega_{c+}z/c} \sum_{n=0}^{\infty} \hat{\sigma}^{(2n+1)}_{13} e^{in[(\omega_{c+}+\omega_{c-})z/c + (\omega_{c+}-\omega_{c-})t]} \tag{7.6}$$
$$+ e^{-i\Delta_- t - i\omega_{c-}z/c} \sum_{n=0}^{-\infty} \hat{\sigma}^{(2n-1)}_{13} e^{in[(\omega_{c+}+\omega_{c-})z/c + (\omega_{c+}-\omega_{c-})t]} \tag{7.7}$$

The coupling between the terms is straightforward, as it depends simply on the cou-

pling of each control field with the relevant spinwave:

$$\partial_t \hat{\sigma}_{13}^{(\pm 1)} = -(\Gamma - i\Delta)\hat{\sigma}_{13}^{(\pm 1)} + ig\hat{\mathcal{E}}_{\pm} + i\Omega_{c\pm}\hat{\sigma}_{12}^{(0)} + i\Omega_{c\mp}\hat{\sigma}_{12}^{(\pm 2)} \tag{7.8}$$

$$\partial_t \hat{\sigma}_{13}^{(\pm(2n-1))} = -(\Gamma_n - i\Delta)\hat{\sigma}_{13}^{(\pm 1)} + i\Omega_{c\pm}\hat{\sigma}_{12}^{(\pm(2n-2))} + i\Omega_{c\mp}\hat{\sigma}_{12}^{(\pm 2n)} \tag{7.9}$$

$$\partial_t \hat{\sigma}_{12}^{(0)} = -(\gamma + i\delta)\hat{\sigma}_{12}^{(0)} + i\Omega_{c+}^{*}\hat{\sigma}_{13}^{(+1)} + i\Omega_{c-}^{*}\hat{\sigma}_{13}^{(-1)} \tag{7.10}$$

$$\partial_t \hat{\sigma}_{12}^{(\pm 2n)} = -(\gamma_n + i\delta)\hat{\sigma}_{12}^{(2n)} + i\Omega_{c+}^{*}\hat{\sigma}_{13}^{(2n+1)} + i\Omega_{c-}^{*}\hat{\sigma}_{13}^{(2n-1)} \tag{7.11}$$

The equations above show how coupling of each coherence to ever higher momentum coherences is possible. This makes it difficult to solve the equations analytically. To simplify the problem the HOCs can be given a decay term (γ_n) that depends on the atomic motion. Due to the rapid spatial variation the higher order coherences decay quickly for moving atoms. This allows truncation of the equations to just a few terms, while discarding all the terms above $n = 0$ is known as the secular approximation.

For example, by truncating at $n = 1$ for the $\hat{\sigma}_{12}$ operators and $n = 0$ for the $\hat{\sigma}_{13}$ operators, the following equations are found:

$$\partial_t \hat{\sigma}_{13}^{(\pm 1)} = -(\Gamma + i\Delta_{\pm})\hat{\sigma}_{13}^{(\pm 1)} + ig\hat{\mathcal{E}}_{\pm} + i\Omega_{c\pm}\hat{\sigma}_{12}^{(0)} + i\Omega_{c\mp}\hat{\sigma}_{12}^{(\pm 1)} \tag{7.12}$$

$$\partial_t \hat{\sigma}_{12}^{(\pm 2)} = -(\gamma_{(2)} + i\delta)\hat{\sigma}_{12}^{(\pm 2)} + i\Omega_{c\mp}^{*}\hat{\sigma}_{13}^{(\pm 1)} \tag{7.13}$$

$$\partial_t \hat{\sigma}_{12}^{(0)} = -(\gamma_{(0)} + i\delta)\hat{\sigma}_{12}^{(0)} + i\Omega_{c+}^{*}\hat{\sigma}_{13}^{(+1)} + i\Omega_{c-}^{*}\hat{\sigma}_{13}^{(-1)} \tag{7.14}$$

$$(\partial_t \pm c\partial_z)\hat{\mathcal{E}}_{\pm} = igN\hat{\sigma}_{13}^{\pm}. \tag{7.15}$$

7.2 EIT stationary light

In the case of EIT the one-photon detunings are small, $\Delta_{\pm} \ll \Gamma$. We renormalise the Maxwell-Bloch equations so they depend on easily measurable experimental parameters by making a few approximations [121]. First, we use the optical depth $d = g^2 NL/(\Gamma c)$, introduce the renormalised collective operators $\hat{P}_{\pm} = \sqrt{n}\hat{\sigma}_{13}^{(\pm 1)}$, $\hat{S}_{\pm} = \sqrt{n}\hat{\sigma}_{12}^{(\pm 2)}$, $\hat{S} = \sqrt{n}\hat{\sigma}_{12}^{(0)}$, normalise the length $\xi = z/L$, and scale the probe field operator $\hat{\mathcal{E}}_{\pm} \to \sqrt{c/(\Gamma L)}\hat{\mathcal{E}}_{\pm}$. We remove the ∂_t term from the optical equation of motion with the argument that time scales of c/L are orders of magnitude smaller than the time scale of the interaction. Then, making the adiabatic approximation $\Gamma\hat{P}_{\pm} \gg \partial_t\hat{P}_{\pm}$ we arrive at

$$\hat{P}_{\pm} = i\sqrt{d}\hat{\mathcal{E}}_{\pm} + \frac{i\Omega_{\pm}}{\Gamma}\hat{S}_0 + \frac{i\Omega_{\mp}}{\Gamma}\hat{S}_{\pm} \tag{7.16}$$

$$\partial_t \hat{S}_0 = -\gamma S_0 + i\Omega_{+}^{*}\hat{P}_{+} + i\Omega_{-}^{*}P_{-} \tag{7.17}$$

$$\partial_t \hat{S}_{\pm} = -\gamma_{\pm}\hat{S}_{\pm} + i\Omega_{\mp}^{*}\hat{P}_{\pm} \tag{7.18}$$

$$\partial_\xi \hat{\mathcal{E}}_{\pm} = i\sqrt{d}\hat{P}_{\pm}. \tag{7.19}$$

Ignoring the coupling to the higher order coherences, these look like the standard EIT equations with an extra counterpropagating probe field. In fact, a compact analytic solution can be found by ignoring $\hat{S}_{\pm}$:

$$\left[\partial_t + \Gamma\tan^2\theta\left(\cos 2\phi\partial_\xi - \frac{1}{d}\partial_\xi^2\right) + \gamma\right]\hat{S}_0 = 0 \tag{7.20}$$

with the quantities

$$\tan^2\theta = \frac{|\Omega|^2}{d\Gamma^2},\ \tan^2\phi = \frac{|\Omega_-|^2}{|\Omega_+|^2},\ \text{and}\ |\Omega|^2 = |\Omega_+|^2 + |\Omega_-|^2. \tag{7.21}$$

This describes a polariton whose motion depends on the relative Rabi frequencies of the two control fields, and whose optical intensity and diffusion depend on the total intensity of the control fields.

Regarding the coupling of the excited state to the hyperfine state, atomic motion in warm atoms generally prevents the coupling between a probe and the counterpropagating control field. This coupling generates a spinwave with much higher momentum, as it involves the absorption of a photon and a re-emission in the opposite direction. These high momentum spinwaves are known as higher-order coherences (HOCs) and have a wavelength of half the optical wavelength or smaller. We use the term 'coherence' as HOCs include both spinwaves (ground state coherences) and excited state coherences. The movement of warm atoms across that distance (under the experimental conditions of [117] and other demonstrations in warm atoms) occurs much more quickly than the EIT inverse bandwidth, and washes out the sub-wavelength spinwave more quickly than the light can couple to it. In very cold atoms, the washing out occurs much more slowly, and these HOCs can affect the propagation of light. This will be discussed further in Sec. 7.2.3.

7.2.1 Multi-wave mixing

The mechanism shown above to produce stationary light is a multi-wave-mixing process where the spinwave is coupled to both forward and backward travelling probe fields by the respective control field. Under control field parameters that give stationary light conditions, the resulting polariton has a zero group velocity. The interpretation for the stationary light in this case is that the polariton is prevented from spreading by the interference between light travelling along multiple different paths. That is, the light is reflected at a continuum of different points in space, generating interference. This is equivalent to the mechanism by which a Bragg (absorption) grating produces a bandgap.

There is therefore not much distinction between the two interpretations, but the treatment based on multi-wave-mixing is necessary to account for the interactions that generate higher order coherences. That treatment above is based on previous work on the multi-wave-mixing description of stationary light.

This work started with Moiseev and Ham exploring the bandgap characteristics, showing that stationary light could be directly generated from a travelling slow light pulse by adiabatically switching on the counterpropagating control field [122]. They further showed this could be useful for wavelength conversion, by using a different wavelength for the backward control field [123], and adiabatically switching off the forward control field after stationary light is formed. The use of a third control field was also examined [124].

In a similar direction, Cui et al. examined the generation of two-colour bandgap and the relationship between the bandgaps [125]. Based on a bandgap derivation, it was also shown that the anomalous dispersion with stationary light could allow for pulse recompression to correct fibre dispersion [126].

A detailed explanation of stationary light based on multi-wave mixing was given by Zimmer et al. [127]. The authors drew upon the pulse matching as described for EIT

[128] and also derived the characteristic spreading time of the stationary light pulse due to the finite optical depth.

The secular approximation can be made exact by ensuring the absence of coupling between the optical coherences and their counterpropagating control fields [129]. The authors proposed a dual-V level scheme to achieve this, illustrated in Fig. 7.2 c). Each optical coherence then only couples to a single control field.

It is worth mentioning that the original stationary light scheme proposed specifically for warm atoms [116], illustrated in Fig. 7.2 b), works differently to the schemes that have since been demonstrated. The probes travelling in opposite directions are coupled via EIT to spinwaves on separate hyperfine states. Since the two probes do not interact with a single spinwave, an extra pair of control fields is used to couple the two hyperfine states. This generates the interference effect that prevents the two probe fields from diverging and travelling outward. The dynamics of this stationary light depends on the relative strengths of the individual EIT schemes and the extra control fields coupling the two hyperfine states.

7.2.2 Phase-matching

The mathematical description of the stationary light in [117] is equivalent to a description of multi-wave mixing in the secular regime. This is not surprising as either mechanism creates a coherent interchange of forward travelling light with backward travelling light, resulting in the same behaviour.

Phase-matching is important for stationary light. It is important that the reflected components be phase-matched through space so that the individual reflections will coherently sum instead of destructively interfere. EIT requires phase-matching even absent stationary light. This is not frequently an issue as normally only one control laser mode is used for storage and retrieval. Where the light is retrieved in the backward direction phase matching becomes important [130]. Depending on the length over which the pulse is stored and the phase-mismatch, the efficiency of recall can be significantly reduced. Similarly, within the stationary light, a significant extra decay can be introduced.

The counterpropagating fields can be phase-matched by tuning the angles of the laser fields. The angles are chosen such that the momentum of the spinwave is equal in amplitude and direction for both probe-control pairs. This is not generally possible in warm vapour experiments. Large angles between the probe-control pairs introduce Doppler broadening of the two-photon line, widening the EIT and allowing absorption of the probe. Depending on the wavelengths of the probe and control, exactly counterpropagating probes and angular phase-matching may be mutually exclusive.

Under EIT, the dispersion of the probe can be modified by adjusting the two-photon detuning [131, 132]. This occurs naturally by pulse matching in stationary light. This method is described in the early work on stationary light [116], and used in [117, 118] which demonstrate that the detuning for the maximal reflection of the probe off an ensemble illuminated by counterpropagating control fields is offset slightly from the centre of the EIT window. This is the frequency at which the phase-matching of the multi-wave mixing is optimal.

Driving of the transitions off two-photon resonance can cause phase variation [123]. This is demonstrated experimentally by [133]. Where the pulse bandwidth exists in a region of nonlinear dispersion, the pulse will suffer some chirp or distortion under EIT [134]. This can place lower bounds on the EIT bandwidth where the two-photon reso-

nance is used to tune the phase-matching, although the effect vanishes at large optical depths. This may be an obstacle to reliable cross-phase modulation if the output phase is too sensitive to the experimental parameters. In the context of EIT stationary light, the phase-matching may be non-uniform due to the bandwidth of the pulse, causing additional decay of the stationary light.

7.2.3 Higher order coherences

Higher order coherences are the rapidly spatially varying coherence terms described above. The generation of HOCs is first mentioned by Moiseev and Ham [122], and is examined thoroughly by Hansen and Mølmer [135, 136]. In particular, the later authors examine how HOCs can negatively affect the generation of stationary light.

The additional interference between light generated from these various HOCs can disturb the multi-wave mixing stationary light effect [136, 137]. The polariton can then diverge and travel through the ensemble. This effect and the suppression of it by using a different EIT bandwidth was shown experimentally [138]. The intensity time integral of the stationary light is shown to be finite, even without ground state dephasing [139]. This is in contrast to earlier work showing that with the correct coupling the effect of the HOCs can counterbalance the diffusion of stationary light normally present in the secular regime [136].

A convincing demonstration of the dependence of stationary light on the multi-wave mixing process was presented in [140], where the forward probe and control wavelengths had a wavelength of 780 nm and the backward pair 795 nm. The travelling wave due to the wavelength difference modulated the control field intensity grating several orders of magnitude more quickly than the EIT bandwidth. The energy difference also prevented the direct coupling between forward probe and backward control and vice versa. However, the large wavelength difference prevented effective phase matching of the counterpropagating fields, causing rapid decay of the spinwave during the stationary light period.

Accounting for the Doppler broadening due to the finite atomic temperature [120] allowed a better explanation of the observations of stationary light performed [140].

In [141], a modified form of stationary light called quasistationary light is proposed, where the spatial spreading of the pulse is greatly suppressed while the average group velocity is nonzero. This is achieved by a mismatch in the two control field strengths.

HOCs can be expanded further for systems where the spatial distribution of emitters in relation to the coherences can be finely adjusted. For example, in [142] the effect of the coupling scheme and the placement of atoms on the dispersion relations for the stationary light is examined. It is shown that the method of truncation of the HOCs and the truncation size for analytically treating the HOCs is particularly important.

7.3 Nonlinearities with stationary light

In Chapter 13, we motivate our study of stationary light by investigating how it may enhance nonlinearities for the purpose of cross-phase modulation via the AC-Stark shift. This involves the interaction of stationary light produced from one state with another state, via the atomic medium. In this case it is important that the presence of each state only affects the evolution of the other state in a simple way. Methods of enhancing cross-

phase modulation in particular is a rapidly developing field which we review in Chapter 12.

In this section, we review self nonlinearities, where the intensity of a light field affects the interaction of that light with the ensemble. This nonlinearity is typically produced by AC-Stark shift between the hyperfine state and a 4th level, as illustrated in Fig. 7.3.

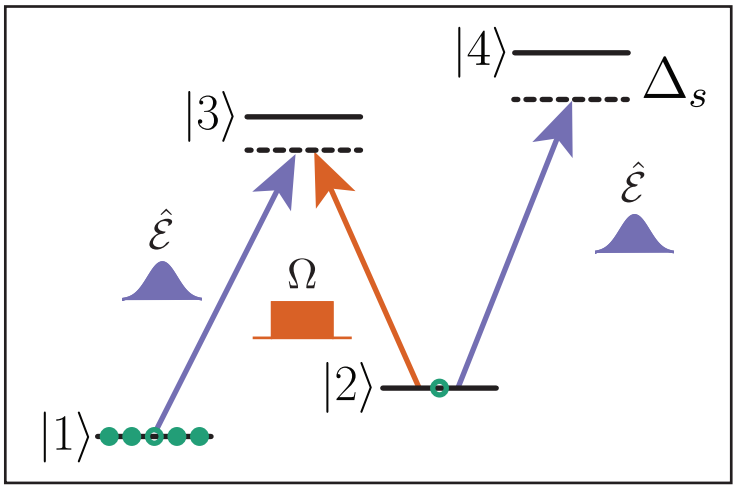

Fig. 7.3. Typical scheme for self nonlinearity. The atoms in state $|2\rangle$ also interact with an additional state $|4\rangle$. The detunings are chosen so $\hat{\mathcal{E}}$ produces an AC-Stark shift of level $|2\rangle$. The intensity of the light therefore modifies the ensemble interaction, generating a nonlinearity. The ensemble interaction can include stationary light.

7.3.1 Photon mass

Under certain nonlinearities, the evolution of the stationary light can be described by models analogous to the Schrödinger equation for massive particles. Making photons mimic the behaviour of massive particles allows the investigation of physical effects that may be more difficult to observe with typical hadronic matter.

For example, it was shown that tightly confining a spinwave in one dimension would allow study of relativistic dispersion at low energies [143]. A similar effect was also proposed in combination with effective magnetic fields for the accurate study of the bosonic quantised Hall effect [144].

7.3.2 Photon crystallisation and condensation

Photon crystals have a regular spatial intensity distribution. This could be converted into a regular time distribution upon readout of the spinwave, allowing for useful time correlations to be generated. Crystallisation may occur due to a two-photon loss term that generates a repulsive force between photons [145, 146]. Alternatively, condensation of the photons can occur if the force is tuned to be attractive [147].

Where the control field forms a standing wave, there is more complicated behaviour. This is due to the interplay between the nonlinearity and the spatial modulation of the interaction generated by the standing wave [148].

7.4 Optical diffraction and phase-matching

The use of HOCs to produce stationary light with a reduced diffusion is mentioned above. However, there are proposals and demonstrations of using standing-wave control fields or spatial gratings in the spinwave to produce useful effects. These do not necessarily involve the use of stationary light, but are based on similar schemes.

For example, there was a proposal to 'blaze' a grating by sending a spatially varying field to modulate the EIT coupling, diffracting the probe [149].

In a proposal for enhancing the probe intensity [150], the probe is sent into an absorption grating generated by standing wave control fields, at the frequency of the first transmission peak (see Fig. 7.4. Bunching or anti-bunching is found depending on the strength of the nonlinearity. The nonlinearity is enhanced due to the increased intensity of light, which reflects many times as it travels through the medium.

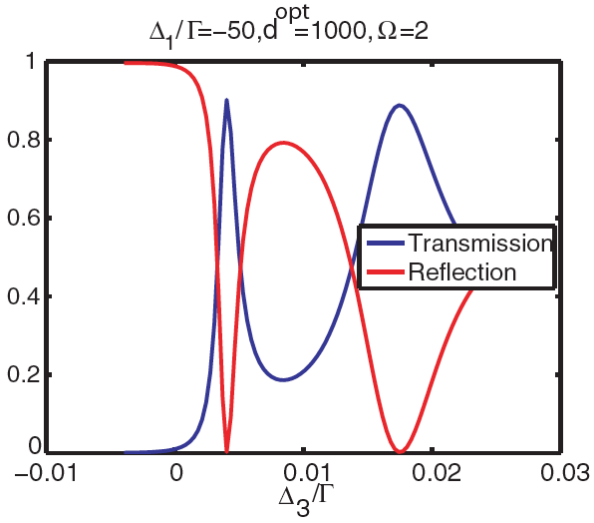

Fig. 7.4. Transmission and reflection from a band gap as a function of detuning. The narrow transmission feature enhances nonlinear interactions at low bandwidth by increasing the intensity of light within the medium. Figure from [151]

A standing wave control field created a spatially modulated level scheme that was used to control the phase-matching of a four-wave mixing signal [152]. The use of four-wave mixing dependent on the standing wave was used to generate an asymmetric diffraction grating [153], see Fig. 7.5. The asymmetry was achieved by generating alternating sections of loss and gain. The diffracted light was emitted only at positive or negative angles, depending on the sign of the dispersion coefficients.

7.5 Stationary light in waveguides

Even with the enhanced interaction due to atomic ensembles and effects such as EIT stationary light, nonlinearities have not been produced to the level necessary for the applications mentioned. In fact, as will be explored later in this thesis, it may be too difficult to achieve this level of nonlinearity in free space. Stationary light might still be combined with optics such as nanofibres or hollow-core photonic fibres to further enhance nonlinear interactions while allowing dynamic control of the interaction. The creation

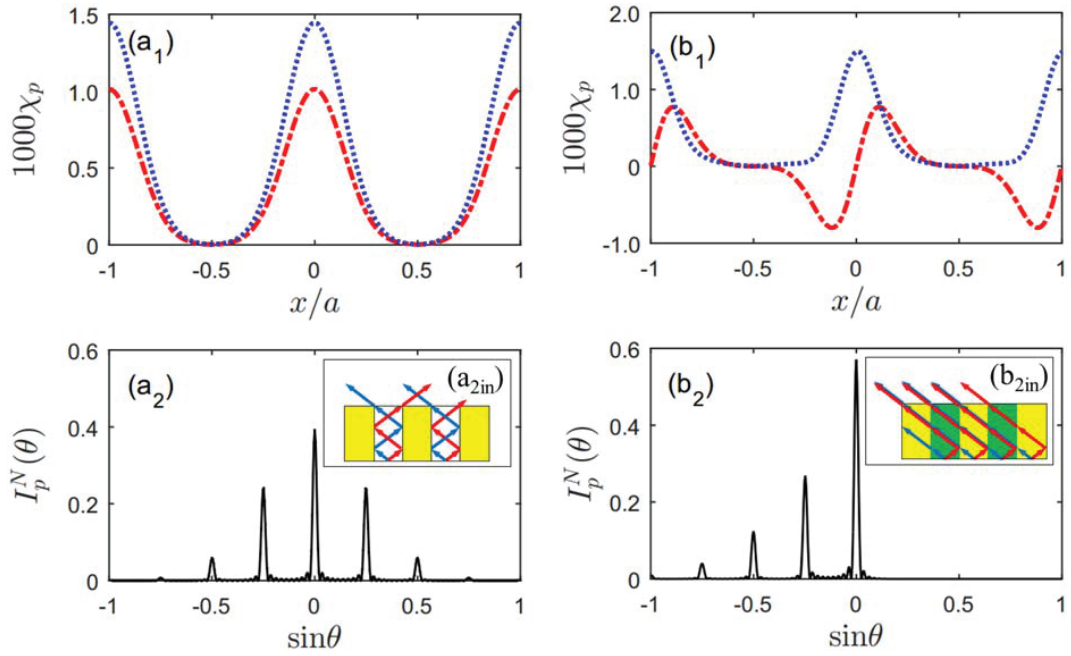

Fig. 7.5. (a_1,b_1) Periodic absorption (red-dashed) and dispersion (blue-dashed), with (a_2,b_2) the corresponding diffraction patterns. Inserts (a_{2in},b_{2in}) show the behaviour of the fields at the loss-gain boundary. Figure from [153]

of stationary light in hollow-core fibre in [154] is a step toward successfully combining these schemes.

7.6 Summary

Stationary light has been explored quite thoroughly in EIT systems. Many new interactions arise due to the presence of counterpropagating fields and the (optional) optical band gap. This review gives the context for my work in the following chapters demonstrating a novel form of stationary light.

Chapter 8

Raman stationary light

This is the flame tree; look how gloriously
That careless blossomer scatters, and more, and more.
What the earth takes of her, it will restore.

Judith Wright, *The Flame Tree*

In this chapter, the theory for a new form of stationary light is derived. This form of stationary light is based on the Raman interaction in which the driving fields are far detuned from the atomic resonance. The literature review in Chapter 7 presented stationary light based on EIT. That form of stationary light relies on a local multi-wave interference mechanism. The propagation of light is prevented due to the interference generated by the multi-wave mixing between the different propagation directions of the light.

By contrast with EIT, the absorption length of light is much greater for large detunings from atomic resonance. The interference between light reflected after travelling short distances can therefore not be relied on to trap the light. Instead, the effect that traps the light in Raman stationary light is an interference effect between light that has not been reflected but rather generated at different positions in the memory. Raman stationary light can potentially be generated over a much larger length.

The relevant published paper to this chapter is

"**Dynamical observations of self-stabilizing stationary light** J. L. Everett, G. T. Campbell, Y.-W. Cho, P. Vernaz-Gris, D. B. Higginbottom, O. Pinel, N. P. Robins, P. K. Lam, B. C. Buchler, Nature Physics 13(1), 68-73 (2016)"

Raman stationary light is created by first sending a probe pulse to write a particular shape of spinwave. The spinwave is then illuminated with counterpropagating control fields. The probe re-emitted throughout the length of the memory interferes destructively by the ends of the memory and no optical field escapes. The spinwaves generated by each probe field travelling in opposite directions within the memory also interfere destructively, cancelling out any evolution within the memory. A bright optical field can then be trapped within the memory.

8.1 Assumptions and approximations

In this chapter, the mathematical description of this behaviour is derived, and the assumptions necessary to arrive at that description are found. We apply the necessary

assumptions to the thorough treatment of the interaction laid out in Chapter 7 to arrive at simple analytic equations for the behaviour of Raman stationary light. We then include some physics that has been left out of the simplified equations in order to more accurately describe the behaviour of this form of stationary light.

We use the Hamiltonian (7.4) as a starting point, and derive the slowly varying operators in a similar way to Sec. 7.1. We make two major assumptions that allow us to write simpler equations of motion that can be solved analytically: the secular approximation, and the phase-matching assumption. We justify these below.

Here is a reminder of the terms involved in the equations: $\hat{\mathcal{E}}_+(z,t)$ and $\hat{\mathcal{E}}_-(z,t)$ are the operators for the probe field amplitudes, with the subscripts $+$ and $-$ indicating travel in the increasing and decreasing z direction respectively. $\Omega_{c+}(t)$ and $\Omega_{c-}(t)$ are the Rabi frequencies for control fields travelling in the same directions. The (z,t) dependence is frequently left out of equations for readability. $\Delta_\pm$ is the detuning of these fields from one-photon resonance. g is the coupling strength of the probe to the atomic transition and N is the linear atomic density.

$\hat{\sigma}_{12}^{\pm}$ are the operators for the ground state atomic coherences, with the superscript $\pm$ referring to the direction of travel of the control-probe pair that generated them, and similarly $\hat{\sigma}_{13}^{\pm}$ are the excited state coherences.

8.1.1 Secular approximation

In calculating the time evolution of the slowly varying operators, we will ignore cross-terms in the Hamiltonian (7.4) of the form $\Omega_{c\pm}\hat{\sigma}_{32}^{\mp}$ and $\hat{\mathcal{E}}_\pm\hat{\sigma}_{31}^{\mp}$. This discards the higher order coherences. Since this theory involves control fields that are far-detuned from each other, we make the assumption that the fastest interaction bandwidth (which is derived below) $d\Gamma\Omega^2/\Delta^2 \ll (\Delta_+ - \Delta_-)$ is much slower than the detuning $\Delta_+ - \Delta_-$. In other words the detuning between the two control fields is much larger than the interaction bandwidth.

8.1.2 Phase matching

The spatial dependency of the phase of each field is important. It is necessary to match the two spinwaves, $\hat{\sigma}_{12}^{+} = \hat{\sigma}_{12}^{-}$ to obtain a single spinwave which each probe field is coupled into and out of. As the two-photon detunings are identical for each field, $\omega_{p+} - \omega_{c+} = \omega_{p-} - \omega_{c-}$, the two spinwave operators $\hat{\sigma}_{12}^{\pm}$ only have a separate spatial dependence given by $\vec{k}_{p\pm} - \vec{k}_{c\pm}$.

$$\partial_t \hat{\sigma}_{12}^{\pm} = i\Omega_{c\pm}^{*}\sigma_{13}^{\pm} e^{i\mathbf{r}\left(\vec{k}_{p\pm} - \vec{k}_{c\pm}\right)}$$

For $|\vec{k}_{c\pm}| > |\vec{k}_{p\pm}|$, it is possible to set $\vec{k}_{p+} - \vec{k}_{c+} \approx \vec{k}_{p-} - \vec{k}_{c-}$ and $\vec{k}_{p+} \parallel \vec{k}_{p-}$ by setting the control fields at an angle to the probe fields. The two probe fields must be parallel so that a one-dimensional description is sufficient. The two momenta cannot be matched exactly as the two control-fields have different frequencies. However, the k-vector mismatch due to the difference in frequency is negligible for any realistic detuning. Detunings up to several GHz are reasonable and the optical frequency is about 400 THz, giving a mismatch of $\Delta_k/k = 10^{-6}$. A typical experiment in cold atoms has a probe beam waist of 100 μm, which is about the wavelength of the spinwave in the transverse dimension. The mismatch would therefore only be important if the spinwave was 10^6 times larger in

the transverse dimension. For non-plane wave fields there is a momentum distribution rather than a single momentum, but this will be addressed in Sec. 8.5.

Once the additional spatial terms have been discarded by using appropriate phase matching, we can set $\hat{\sigma}_{12} = \hat{\sigma}_{12}^{+} + \hat{\sigma}_{12}^{-}$.

The familiar co-propagating three-level equations then appear, with an extra probe and corresponding optical coherence all interacting with a single spinwave:

$$\begin{aligned}
\partial_t \hat{\sigma}_{13}^{+} &= -(\Gamma + i\Delta_{+})\hat{\sigma}_{13}^{+} + ig\hat{\mathcal{E}}_{+} + i\Omega_{c+}\hat{\sigma}_{12} \\
\partial_t \hat{\sigma}_{13}^{-} &= -(\Gamma + i\Delta_{-})\hat{\sigma}_{13}^{-} + ig\hat{\mathcal{E}}_{-} + i\Omega_{c-}\hat{\sigma}_{12} \\
\partial_t \hat{\sigma}_{12} &= -(\gamma + i\delta)\hat{\sigma}_{12} + i\Omega_{c+}^{*}\hat{\sigma}_{13}^{+} + i\Omega_{c-}^{*}\hat{\sigma}_{13}^{-} \\
(\partial_t + c\partial_z)\hat{\mathcal{E}}_{+} &= igN\hat{\sigma}_{13}^{+} \\
(\partial_t - c\partial_z)\hat{\mathcal{E}}_{-} &= igN\hat{\sigma}_{13}^{-}.
\end{aligned} \tag{8.1}$$

The equations can be simplified by removing the optical coherences, $\hat{\sigma}_{13}^{\pm}$. The standard adiabatic approximation is used as in (2.73). Unlike in the single directional case, the terms describing the dispersion of the probe fields ($\partial_z \hat{\mathcal{E}}_{+} \rightarrow ig^2 N\hat{\mathcal{E}}_{+}/(c\Delta_{+})$) should be considered.

$$\partial_t \hat{\sigma}_{12} = -(\gamma + i\left(\delta - \frac{|\Omega_{c+}|^2}{\Delta_{+}} - \frac{|\Omega_{c-}|^2}{\Delta_{-}}\right))\hat{\sigma}_{12} + ig\left(\frac{\Omega_{c+}^{*}}{\Delta_{+}}\hat{\mathcal{E}}_{+} + \frac{\Omega_{c-}^{*}}{\Delta_{-}}\hat{\mathcal{E}}_{-}\right) \tag{8.2}$$

$$(\partial_t + c\partial_z)\hat{\mathcal{E}}_{+} = i\frac{gN\Omega_{c+}}{\Delta_{+}}\hat{\sigma}_{12} + i\frac{g^2 N}{c\Delta_{+}}\hat{\mathcal{E}}_{+} \tag{8.3}$$

$$(\partial_t - c\partial_z)\hat{\mathcal{E}}_{-} = i\frac{gN\Omega_{c-}}{\Delta_{-}}\hat{\sigma}_{12} + i\frac{g^2 N}{c\Delta_{-}}\hat{\mathcal{E}}_{-}. \tag{8.4}$$

The equations can be simplified further by writing them in terms of more convenient parameters. The length over which z runs is normalised to $\xi = z/L$. and the optical depth, $d = g^2 N/(\Gamma c)$ is substituted in to remove linear atomic density N and interaction strength g. To completely replace terms containing the linear atomic density, the coherences are also renormalised to this density: $S = \sqrt{N}\hat{\sigma}_{12}$ and $P_{\pm} = \sqrt{N}\hat{\sigma}_{13}^{\pm}$. The probe field is also renormalised $\hat{\mathcal{E}}_{\pm} \rightarrow \sqrt{c/\Gamma}\hat{\mathcal{E}}_{\pm}$. The optical depth is a standard parameter determining the efficiency of optical memories, and we show later that it is also important in determining the amount of useful nonlinearity possible with Raman stationary light.

The time derivatives ∂_t of the probe fields are neglected. This is justified by the rapid travel time of the probe through the memory, which for a 3 cm memory is $L/c = 0.1$ ns. For comparison, the memory evolves on the time scale of microseconds. We then arrive at the following equations:

$$\partial_t \hat{S} = -(\gamma + i\left(\delta - \frac{|\Omega_{c+}|^2}{\Delta_{+}} - \frac{|\Omega_{c-}|^2}{\Delta_{-}}\right)\hat{S} + i\sqrt{d}\Gamma\left(\frac{\Omega_{c+}^{*}}{\Delta_{+}}\hat{\mathcal{E}}_{+} + \frac{\Omega_{c-}^{*}}{\Delta_{-}}\hat{\mathcal{E}}_{-}\right) \tag{8.5}$$

$$\partial_\xi \hat{\mathcal{E}}_{+} = i\sqrt{d}\frac{\Omega_{c+}}{\Delta_{+}}\hat{S} + i\frac{\Gamma d}{\Delta_{+}}\hat{\mathcal{E}}_{+} \tag{8.6}$$

$$\partial_\xi \hat{\mathcal{E}}_{-} = -i\sqrt{d}\frac{\Omega_{c-}}{\Delta_{-}}\hat{S} - i\frac{\Gamma d}{\Delta_{-}}\hat{\mathcal{E}}_{-}. \tag{8.7}$$

8.2 The stationary light equation

These Eqs. (8.5)-(8.7) can be simplified further by setting the control field Rabi frequencies equal, $\Omega_{c+} = \Omega_{c-} = \Omega$, and their detunings from the excited state equal and opposite $\Delta_+ = -\Delta_- = \Delta$. We also set each pair of fields on two-photon resonance, $\delta = 0$. Both probe equations of motion now have the same dispersion term $i\Gamma d/\Delta$. By transforming into a rotating spatial frame

$$\hat{\mathcal{E}}_\pm \rightarrow \hat{\mathcal{E}}_\pm e^{(i\frac{\Gamma d}{\Delta}\xi)}, \tag{8.8}$$

and similarly for the spinwave, this dispersion term can be removed:

$$\partial_t \hat{S}(t,\xi) = i\sqrt{d}\,\frac{\Gamma\Omega}{\Delta}\left(\hat{\mathcal{E}}_+ + \hat{\mathcal{E}}_-\right) - \gamma\hat{S} \tag{8.9}$$

$$\partial_\xi \hat{\mathcal{E}}_+(t,\xi) = i\sqrt{d}\frac{\Omega}{\Delta}\hat{S} \tag{8.10}$$

$$\partial_\xi \hat{\mathcal{E}}_-(t,\xi) = -i\sqrt{d}\frac{\Omega}{\Delta}\hat{S}. \tag{8.11}$$

The evolution of the spinwave then has a simple analytic solution. The probe fields throughout the memory are found by integrating over the spinwave. Where no probes are input into the memory, the boundary condition is zero at the input end for that probe. This gives the probe fields,

$$\hat{\mathcal{E}}_+(t,\xi) = i\sqrt{d}\frac{\Omega}{\Delta}\int_0^\xi \hat{S}(t,\xi')\mathrm{d}\xi' \tag{8.12}$$

$$\hat{\mathcal{E}}_-(t,\xi) = -i\sqrt{d}\frac{\Omega}{\Delta}\int_1^\xi \hat{S}(t,\xi')\mathrm{d}\xi'. \tag{8.13}$$

Substituting these expressions for the probe field into (8.9), we arrive at the stationary light equation:

$$\begin{aligned}(\partial_t + \gamma)\,\hat{S}(t,\xi) &= -d\,\Gamma\frac{\Omega^2}{\Delta^2}\left(\int_0^\xi \hat{S}(t,\xi')\mathrm{d}\xi' - \int_1^\xi \hat{S}(t,\xi')\mathrm{d}\xi'\right)\\ &= -d\,\Gamma\frac{\Omega^2}{\Delta^2}\int_0^1 \hat{S}(t,\xi')\mathrm{d}\xi'.\end{aligned} \tag{8.14}$$

Equation (8.14) states that, apart from the decay term γ, any spinwave that integrates to zero along the length of the memory will not evolve. According to equations (8.12), (8.13), no probe light is output from such a spinwave. A spinwave that integrates to zero along the length of the memory will therefore be referred to as a stationary spinwave, and the light generated under the appropriate control fields as stationary light. The assumptions made to reduce the Maxwell-Bloch equations to this form are necessary to arrive at this analytic solution. However, we will investigate later in this chapter how the stationary light behaves when these assumptions are relaxed.

The behaviour of spinwaves that do not satisfy this solution can still be described by the stationary light equation. A constant term can always be subtracted from any spinwave so that it integrates to zero, and becomes a stationary spinwave. Any spinwave can then be written as the sum of a stationary and non-stationary component $\hat{S}(\xi,t) =$

$\hat{S}_\xi(\xi) + \hat{S}_t(t)$ where $\int_0^1 \hat{S}_\xi(\xi')\mathrm{d}\xi' = 0$. Remembering that the evolution of the spinwave is independent of position, the stationary light equation can be applied at time $t = 0$.

$$(\partial t + \gamma)\hat{S}(t,\xi) = -d\frac{\Gamma\Omega^2}{\Delta^2}\int_0^1 \hat{S}(t,\xi')\mathrm{d}\xi' \tag{8.15}$$

$$= -d\frac{\Gamma\Omega^2}{\Delta^2}\hat{S}_t(t) \tag{8.16}$$

This is a uniform evolution in space, even including the decay term, so the spinwave can now be described by the equation

$$\hat{S}(\xi,t) = [\hat{S}_\xi(\xi) + \hat{S}_t(t)]e^{-\gamma t} \tag{8.17}$$

Applying the evolution equation to the spinwave gives the time evolution of the (spatially) constant term $\hat{S}_t(t) = \hat{S}_{t_0}.\exp\left(-t(d\Gamma\Omega^2/\Delta^2)\right)$. This means that any spinwave will evolve to a stationary spinwave. Where the spinwave contains no stationary component and is constant in space, it will evolve to zero, and all the light will exit the memory. This is described by the equation

$$\hat{S}(\xi,t) = [\hat{S}_\xi(\xi) + \hat{S}_{t_0}e^{-t(\frac{d\Gamma\Omega^2}{\Delta^2})}]e^{-\gamma t} \tag{8.18}$$

The intuitive explanation for this behaviour can be found by considering Eq. (2.76). The driving of the spinwave is proportional to the sum of the two probe fields. Equations (8.10), (8.11) and the fact that the optical inputs are zero at each end mean that a spinwave integrating to zero also give zero optical outputs. Then, integrating the optical fields from either end using Eqs. (8.10), (8.11), it can be seen that the optical fields are equal and opposite to one another throughout the memory.

If the spinwave is not initially stationary, the sum of the two counterpropagating optical fields is still constant along the z dimension. The driving of the spinwave is then constant with z. The spinwave then evolves until the integral is zero, with the spatially constant spinwave term decaying due to the spatially constant driving term. The spinwave becomes stationary, while the optical fields are non-zero within the memory but zero at the ends, and stationary light is produced.

This behaviour can be neatly illustrated by ignoring the incoherent decay of the spinwave. For spinwaves that are entirely real in the stationary light equation frame (no imaginary components), no imaginary terms are produced and the evolution can be illustrated in a 2-dimensional plot. Using two similar spinwaves, the different behaviours of initially stationary and non-stationary spinwaves are described.

We consider a spinwave composed of two spatially separated Gaussians. In the case where one Gaussian is negative and one positive (the phase between the Gaussians is π) the spinwave will integrate to zero along the length of the memory. In the case where the two Gaussians have the same phase, the spinwave will not integrate to zero. Fig. 8.1 shows the probe fields produced, their relative phases, and how the two spinwaves behave under stationary light.

The spinwave that integrates to zero is stationary and does not evolve. The optical fields are the integral of the spinwave in their direction of travel. This gives the equal and opposite fields, with the cancellation of the driving of the spinwave.

For the spinwave that does not integrate to zero, the optical fields are non-zero at the ends and the light escapes the memory. The driving of the spinwave no longer cancels

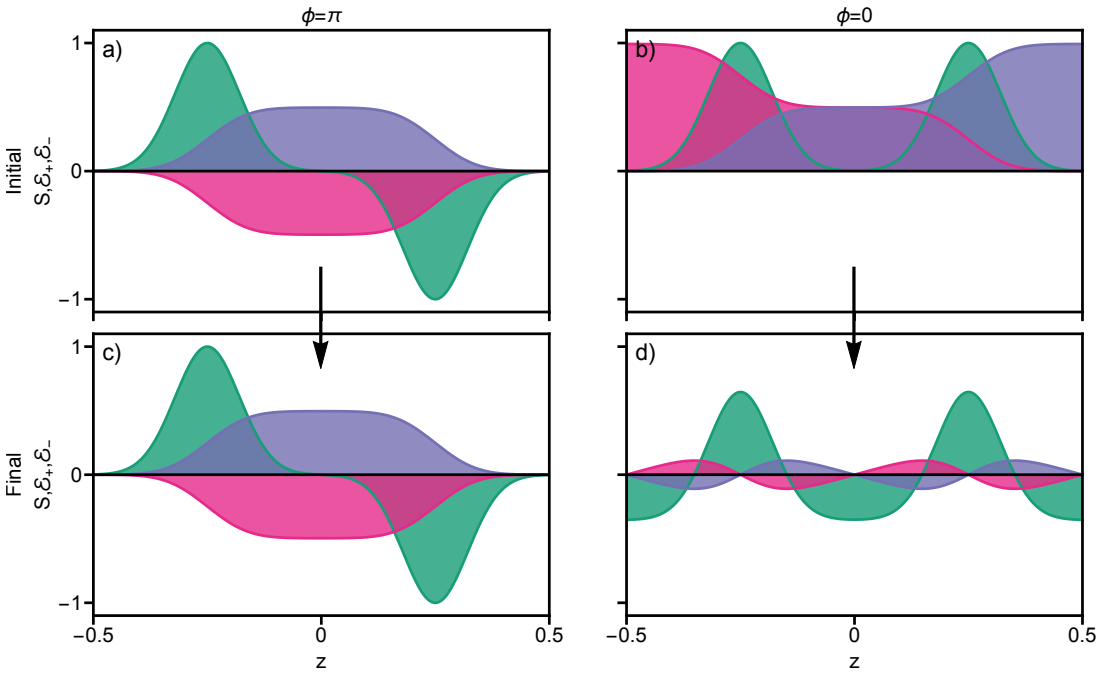

Fig. 8.1. Simplified dynamics of stationary light for two spinwaves evolving under stationary light. Spinwave (green), forward probe $\mathcal{E}_+$ (blue) and backward probe $\mathcal{E}_-$ (purple). a) The initial spinwave with two regions π out of phase integrates to zero along the length and does not evolve. b) The initial spinwave with two regions in phase does not integrate to zero and evolves to d), where the spinwave integrates to zero. The sums of the probes are constant along the length of the memory at all times and so the evolution adds a spatially constant term to the spinwave.

out, and the spinwave evolves to a stationary spinwave. The amplitudes of the final stationary optical fields are much smaller than the fields in the initially stationary case; a large proportion of the initial spinwave has left the memory, and the alternation of opposite phase regions of spinwave reduces the maximum of the integral describing the probe light.

These spinwaves are straightforward to write using the gradient echo memory, and form the basis of the experimental demonstration of stationary light in Chapter 9.

8.3 Simulation of stationary light

The analytic solution for the behaviour of Raman stationary light is based on simplified equations. Modelling the behaviour of the stationary light under more realistic equations will assist with the design and interpretation of an experimental demonstration, and allow for the testing of the decay terms described below. The numerical PDE solver XMDS2 [155] is used to simulate the three-level equations (8.1). The ∂_t terms are removed from the optical equations of motion as with the analytic stationary light solution, as this cannot be simulated efficiently.

Simulations are done with at least 100 lattice points in z or ξ, and more often 200 or 400. Simulations work by calculating the expected optical fields at each lattice point at a point in time based on the partial differential equation, using the Runge Kutta method (RK4). These outputs are then fed into a second solver which calculates the evolution of the spinwave and optical coherences based on the optical fields, using an adaptive Runge Kutta method (ARK45 or ARK89). The adaptive method adjusts the step size of the integration so that the estimated truncation error is less than 10^{-9}. This results in a large number of steps but the output is only sampled at about 200-400 time samples per simulation. These options are all implemented at a high level in the XMDS2 script, which is then compiled and optimised for C by the XMDS2 package.

Example code is included in appendix A.

8.4 Decay of stationary light

Some interactions were neglected in deriving the stationary light equation so that an intuitive form could be found. These interactions can cause additional evolution of the spinwave, and so will need to be considered for practical uses of Raman stationary light. In Chapter 13, by including the unavoidable scattering of the spinwave by the control field, the amount of cross-phase modulation that can be generated with stationary light is shown to depend on the optical depth. We can then explore how any other evolution processes scale with optical depth, to find the regime where these processes do not adversely affect cross-phase modulation or other desired nonlinear interaction.

We define the characteristic decay rate

$$\Gamma_{SL} = d\Gamma\frac{\Omega^2}{\Delta^2}. \tag{8.19}$$

We show in Chapter 13 that the quantity Γ_{SL} is proportional to the intensity of the light. Therefore when multiplied by time it is proportional to the amount of nonlinear interaction generated. It is therefore useful for measuring the decay generated by some of the various mechanisms as it directly scales the decay to the intensity of the stationary

light. It then serves as a comparison point for the different decay mechanisms to see how they will affect the possible nonlinear interaction.

8.4.1 Incoherent absorption of the control fields

The control field has a chance to incoherently scatter off atoms in the hyperfine state. This affects the whole spinwave equally as an exponential loss term, so is not interesting to simulate. The loss rate due to this mechanism is $\Gamma_\Omega = 2\Gamma\Omega^2/\Delta^2$. This is proportional to Γ_{SL}/d, suggesting that the optical intensity of the stationary light integrated over time is proportional to d. We show that this is indeed a limiting factor for nonlinear interactions in Chapter 13.

8.4.2 Incoherent absorption of the probe fields

For the following, we assume we start with a stationary solution of the simplified equations; that the integral of the spinwave $\int_0^1 d\xi.\hat{S} = 0$. We continue to assume that the memory is short enough that any evolution occurs with characteristic time much longer than the travel time of light across the memory. In this case, scattering due to the control fields is a homogeneous decay process and simply reduces the magnitude of the spinwave.

Incoherent absorption of the probe fields causes a more complicated evolution. Under the simplified equations for a stationary spinwave, all of the probe field emitted that is emitted by the memory adds to zero at the ends. In reality, the probe is also incoherently scattered by the atoms and so its amplitude decreases as it travels through the memory according to $\partial_\xi \hat{\mathcal{E}}_\pm = \mp d\frac{\Gamma^2}{\Delta^2}\hat{\mathcal{E}}_\pm$. Integrating over the memory gives

$$\hat{\mathcal{E}}_+(\xi) + \hat{\mathcal{E}}_-(\xi) = i\sqrt{d}\frac{\Omega}{\Delta}\left[\int_0^\xi d\xi'.\hat{S}(\xi')\exp\left(-d\frac{\Gamma^2}{\Delta^2}(\xi-\xi')\right)\right. \tag{8.20}$$

$$\left. + \int_1^\xi d\xi'.\hat{S}(\xi')\exp\left(-d\frac{\Gamma^2}{\Delta^2}(\xi'-\xi)\right)\right]. \tag{8.21}$$

Where the scattering is small, it can be approximated as a linear decay of the probe,

$$\hat{\mathcal{E}}_+(\xi) + \hat{\mathcal{E}}_-(\xi) = i\sqrt{d}\frac{\Omega}{\Delta}\left[\int_0^\xi d\xi'.\hat{S}(\xi')(1-d\frac{\Gamma^2}{\Delta^2}(\xi-\xi'))\right. \tag{8.22}$$

$$\left. + \int_1^\xi d\xi'.\hat{S}(\xi')(1-d\frac{\Gamma^2}{\Delta^2}(\xi'-\xi))\right]. \tag{8.23}$$

The constant term of the spinwave integrates to zero, and

$$\hat{\mathcal{E}}_+(\xi) + \hat{\mathcal{E}}_-(\xi) = id\sqrt{d}\frac{\Omega\Gamma^2}{\Delta^3}\left[\int_0^\xi d\xi'.\hat{S}(\xi')(\xi-\xi') + \int_1^\xi d\xi'.\hat{S}(\xi')(\xi'-\xi))\right]. \tag{8.24}$$

or in another form, indicating how the imbalance between the probe fields depends on the spatial distribution of the spinwave,

$$\frac{d^2}{d\xi^2}\left(\hat{\mathcal{E}}_+(\xi) + \hat{\mathcal{E}}_-(\xi)\right) = 2id\sqrt{d}\frac{\Omega\Gamma^2}{\Delta^3}\hat{S}(\xi). \tag{8.25}$$

Then, the evolution of the optical field caused by incoherent absorption is $\frac{\partial\hat{\mathcal{E}}}{\partial t} =$

$\frac{d^2\Gamma^3\Omega^2}{\Delta^4}\hat{\mathcal{E}}$, which is proportional to $d\frac{\Gamma^2}{\Delta^2}\Gamma_{SL}$. This relationship is important when designing an experiment: If optical depth can be increased in order to increase stationary light intensity, the one-photon detuning Δ should be increased by at least the same proportion to so that loss due to incoherent absorption is not increased.

Simulation data demonstrating the effect of the spinwave distribution on the decay are shown in Fig. 8.2. All the spinwaves correspond to a stored optical pulse of equal intensity, in a memory of the same optical depth. However the distribution of the spinwaves over that optical depth differs. The spinwaves spread over more optical depth decay more quickly, as predicted by the above equation. For example, the spinwaves in Fig. 8.2 a) and d) are identical apart from the optical depth between the regions of spinwave. The stationary light intensity decays much more rapidly in a)-c).

The spinwave in Fig. 8.2 g), has a different shape, and is spread out over almost the entire optical depth. This produces a larger initial stationary light intensity, and a larger intensity integrated over time, than the other spinwaves. The scaling of stationary light intensity with d and decay rate with d^2 indicates there is an ideal optical depth and distribution over which to store a spinwave to produce maximum integrated stationary light intensity, for a given control field detuning Δ.

For single-colour stationary light (where $\Delta_+ = \Delta_-$), dispersion causes the phases of the two probe fields to become unequal. The decay of the SL field due to this mechanism scales more poorly than for two-colour SL, as the dependence of the phase of the probe field is given by the real part of the dispersion, from $\partial_\xi\hat{\mathcal{E}}_\pm = \pm id\frac{\Gamma}{\Delta}\hat{\mathcal{E}}_\pm$. The decay is proportional to $d^2\frac{\Gamma}{\Delta}\Gamma_{SL}$, meaning Δ should be increased proportionally with d^2 to maintain the same loss rate.

Spatial evolution of the spinwave due to incoherent probe absorption is not expected for single-photon states. Instead, the incoherent absorption of a single photon destroys the state altogether. This scattering rate scales in the same manner with optical depth and detuning. This behaviour cannot be analysed in the semi-classical simulations, but the scattering may still affect the behaviour of stationary light beyond simply destroying or not destroying the state.

To maximise the amount of cross-phase shift, the integrated intensity per photon should be maximised. Since intensity and control-field scattering, along with other decay mechanisms above, all scale at the same order with control field power, control field intensity should be maximised so that other non-scaling decays become negligible.

8.4.3 Finite speed of light

To find the stationary light condition in Sec. 8.2, it is assumed that the travel time of light across the memory is negligible compared to the time scale of the evolution of the memory. It is possible to discard this assumption and recover a finite speed of light within simulations. Example code for this simulation is included in appendix A.

Fig. 8.3 displays the results of these simulations with various driving strengths. In a)-c) the evolution is not noticeably different to a simulation where the speed of light is effectively infinite. In d)-f) the spinwave evolves during the time before both control fields cover the entire spinwave. When the generated optical fields reach the other side of the spinwave, they are mostly reabsorbed, and a stationary light field is formed. In g)-i), the spinwave is almost entirely converted to light in the time it takes the control fields to traverse the space between the spinwaves.

Even at the lowest driving strength simulated, the magnitude of the driving term

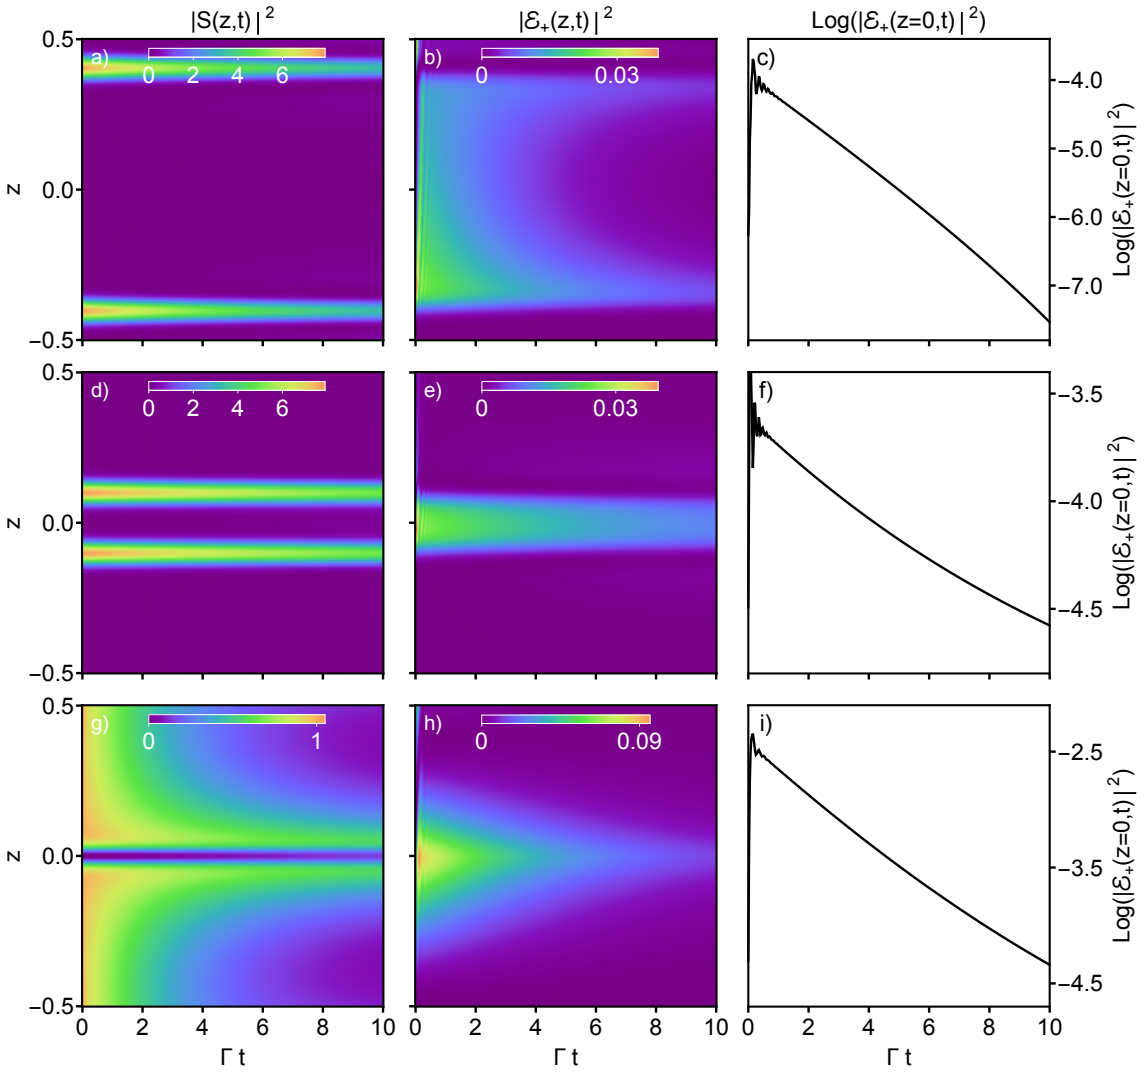

Fig. 8.2. 3-level simulation data comparing incoherent absorption-induced decay for a range of spinwaves. All the spinwaves (a,d,g) are stationary ($\int S \mathrm{d}\xi = 0$) and satisfy $\int |S|^2 \mathrm{d}\xi = 1$. The parameters are $\Omega_c = \Gamma/2$, $\Delta = 10\Gamma$, and $d = 200$. The decay due to incoherent absorption dominates for all the spinwaves. c),f),i) show the decay of the stationary light intensity at the centre of the memory. The decay is not strictly exponential due to the redistribution of the spinwave.

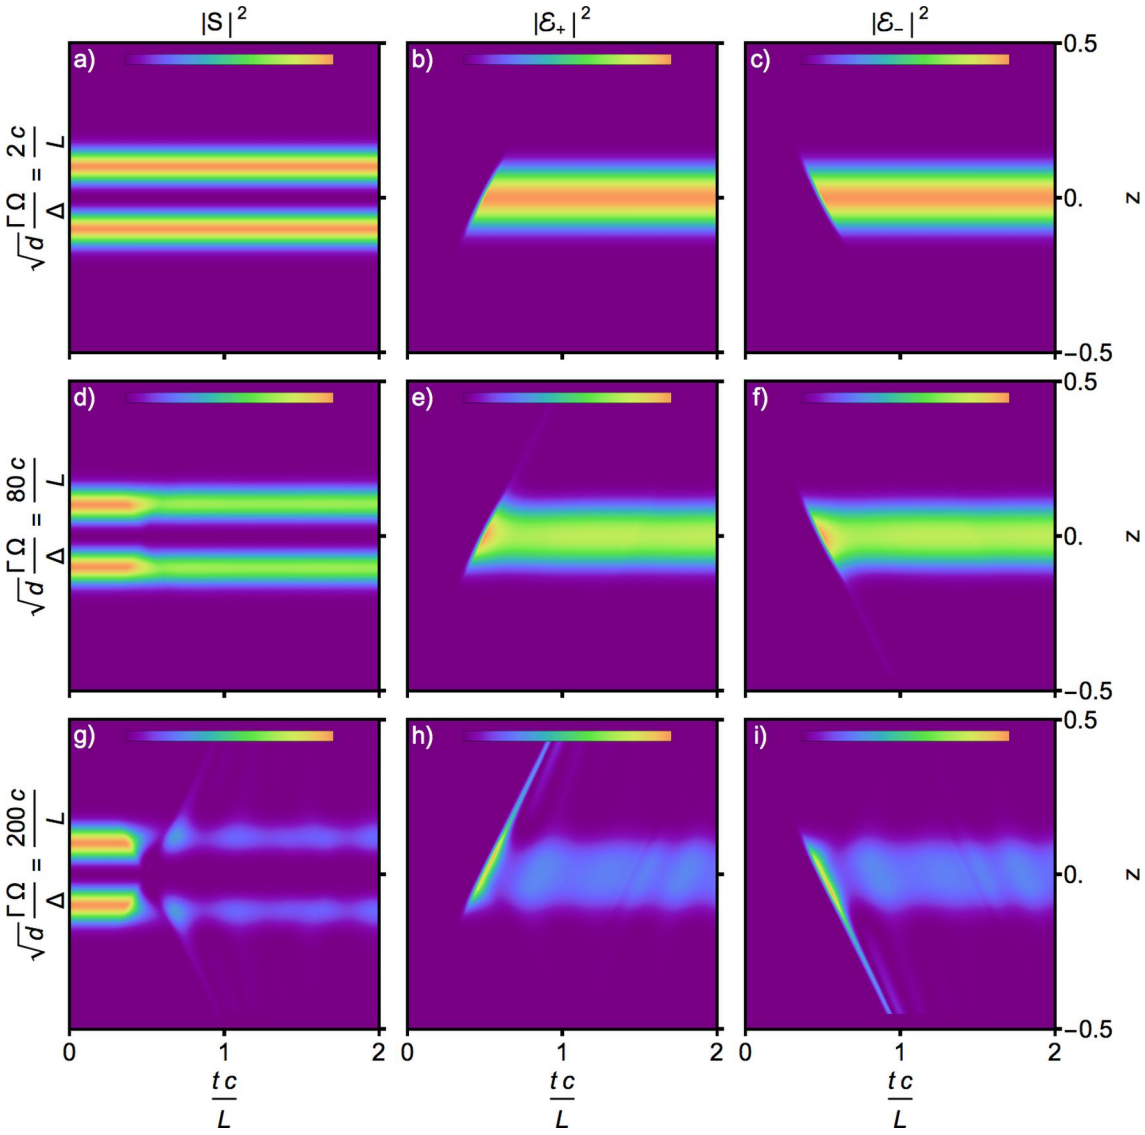

Fig. 8.3. Simulation data for stationary light for a finite speed of light with various driving strengths, and with no decay terms. The control fields enter from either side at the beginning of the simulation. The optical field intensities are not to scale across different driving strengths.

corresponds to about 1 GHz for a memory about 1cm long. For comparison, a typical experimental value would be about 10 MHz. It is hard to imagine a scenario where the larger levels of driving would be possible, particularly without breaking other theoretical assumptions about adiabaticity. Apart from investigating the driving strength at which the stationary light would break down, the motivation for running these simulations was to test if there was any positive feedback due to a finite speed of light. This could amplify the initial single control field driven evolution and destabilise the stationary light. This does not appear to be a problem. In Fig. 8.3 (d) the spinwave clearly stabilises, while even in (g) where the spinwave is mostly erased during the initial evolution, the oscillations appear to be dying down.

The propagation of light that has left the region of the initial spinwave appears to be uni-directional. There is little to no coupling back into the optical field travelling in the opposite direction. This is visible even in Fig. 8.3 h) and i) in the central region. In fact, the propagation of the escaped light out of the memory without coupling back into the counterpropagating light would indicate that there should be no destructive feedback loop in the perturbation of the spinwave.

8.4.4 Control field noise

Noise in the control field will contribute to the decay of the stationary light state. Noise imbalances the two control fields and so the recalled optical fields, driving some evolution of the spinwave. The spinwave then does not integrate to zero and will undergo rapid exponential decay. The magnitude of the decay depends primarily on the amount and type of noise in the control fields, as well as on the distribution of the spinwave. This did not appear to be an issue in the experimental work in Chapter 9, and can be left as a problem for future study.

There will inevitably be some amount of noise in the control fields produced by shot noise. This is negligible in a free space experiment.

8.5 Two and three dimensional stationary light

In the one dimensional study above it is assumed that the probe beams can be perfectly phase-matched. However, in physical two- or three-dimensional optical fields, the fields are no longer plane waves and so will have a transverse momentum distribution. The phase-matching is only perfect for a single value of the momentum, so we will investigate how phase-matching affects stationary light for Gaussian beam profiles.

8.5.1 Diffusion of Gaussian beams

For Gaussian beams, the phase mismatch can be understood by considering the Gouy phase. When a forward propagating beam is imprinted as a spinwave, the Gouy phase, which varies as the beam traverses the waist, will also be imprinted. When the backward propagating beam is generated, the same Gouy phase will be produced, but increasing in the opposite direction. The two beams will no longer be phase matched and will not produce stationary light. This behaviour is illustrated in Fig. 8.4, where a two-dimensional simulation fails to produce completely stationary light, under conditions that would give stationary light in one dimension.

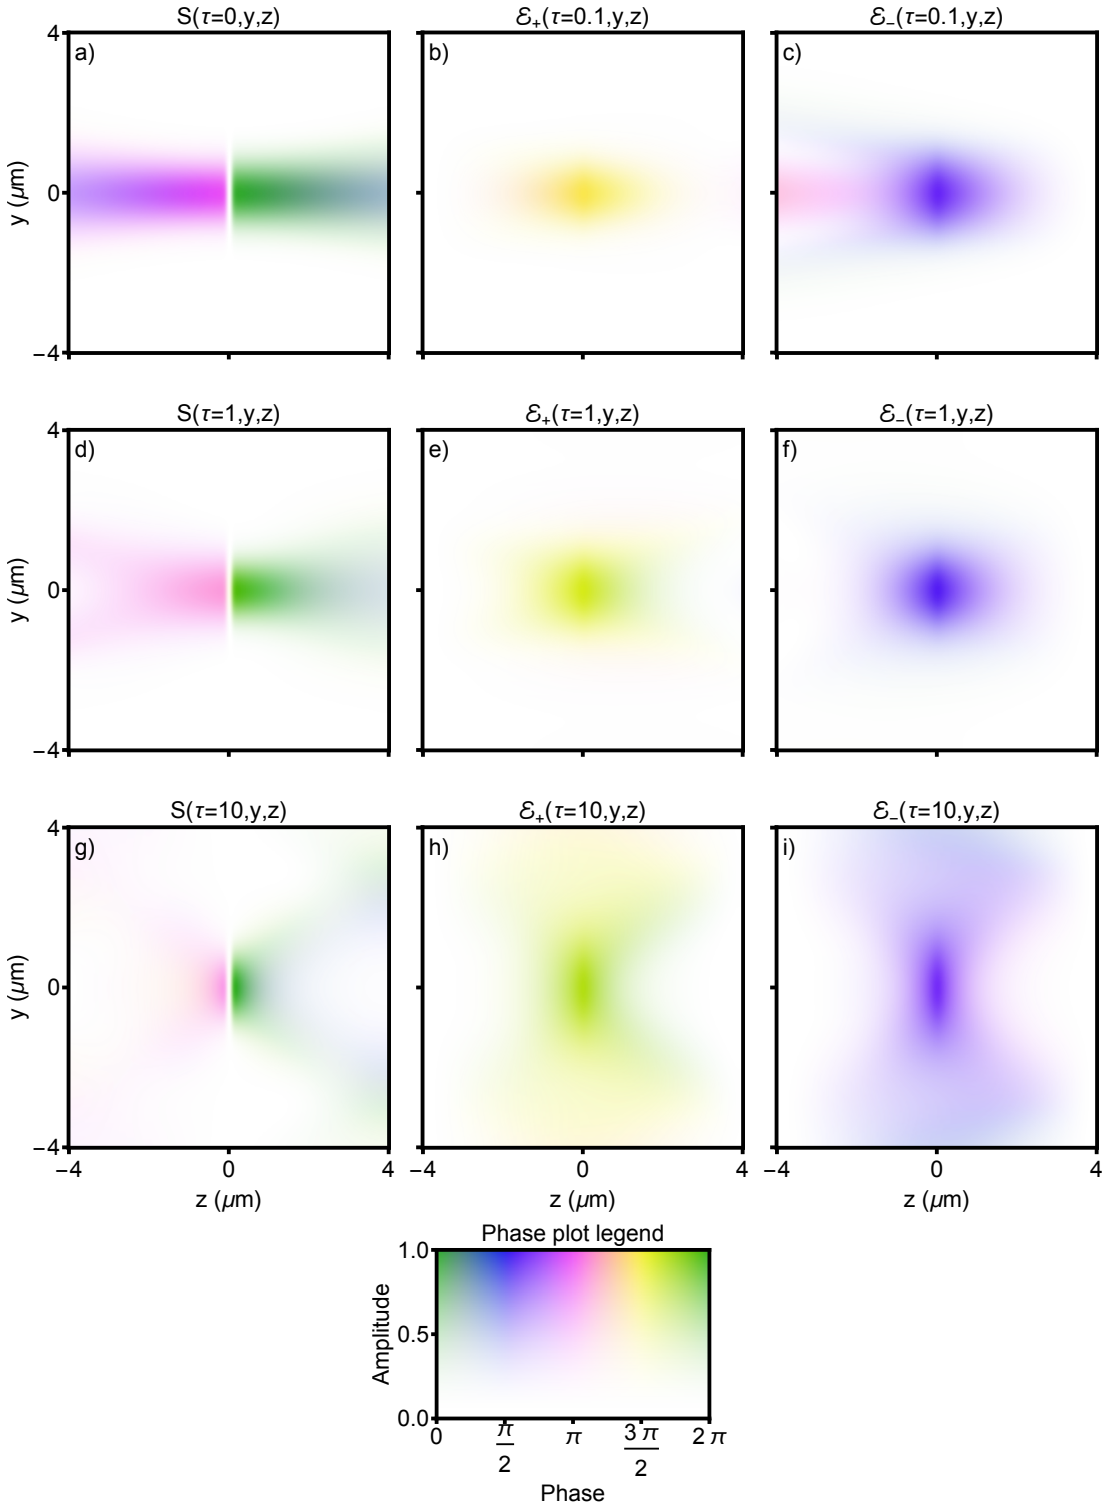

Fig. 8.4. Evolution of a stationary spinwave in a 2d simulation. The stored spinwave corresponds to a Gaussian forward probe stored in an ensemble extending over two Rayleigh lengths. a)-c) show the initial spinwave and fields. b) The initial forward probe is phase-matched throughout the memory while c) the backward probe is not and some escapes. After some time the counterpropagating probes look more like each other. The stationary light continues to decay and diffuse d)-f); there is no self-stabilisation or confinement in two or three dimensions. The amplitude is not to scale between plots; see Fig. 8.5 for the quantitative evolution.

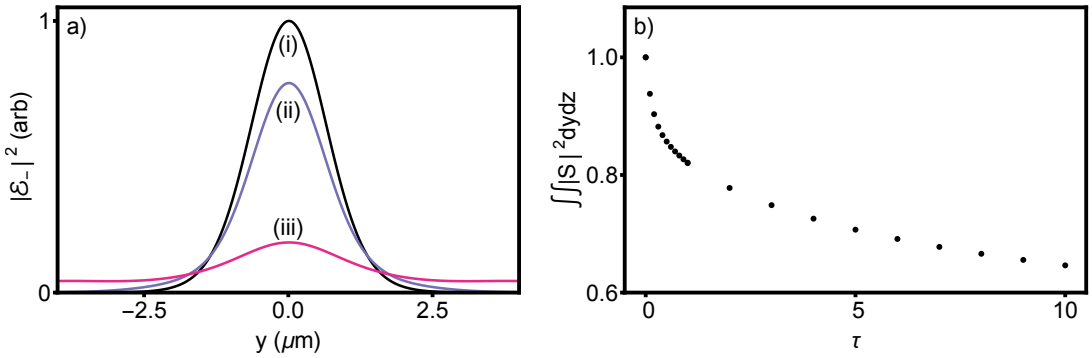

Fig. 8.5. Behaviour of the probe fields and spinwaves for 2D simulation. (a) Intensity $|\mathcal{E}_-|^2$ at $z = 0$ for (i-iii) $\Gamma_{SL}t = (0.1, 1, 10)$. (b) Integrated spinwave remaining compared to initial spinwave.

Due to the constraints of running simulations in higher dimensions while preserving the wavelength-scale features associated with highly focussed beams, a systematic approach to simulating mode preservation is not possible. However, a simple relationship is observed between the number of Rayleigh lengths over which the spinwave is stored, and the speed at which the stationary light decays. This is where all other parameters, including total optical depth, are equal. This can be reconciled with the fact that the quantity $(1-\cos(\theta_{DIV}))L_R$, where L_R is the Rayleigh length, is roughly constant regardless of the waist size of the beam. This factor is proportional to the average phase mismatch of the field that builds up over a Rayleigh length. At short distances this is due to the Gouy phase, while at long distances the phase mismatch is due to the spherical wavefronts of the diverging beam.

The Gouy phase is roughly linear at distances under a Rayleigh length. Therefore it is possible to achieve phase-matching by simply adjusting the control field angle. This does not completely solve the problem, as the beams still diverge. This is also not very satisfactory, as one of the touted advantages of Raman stationary light is the ability to use the full optical depth of ensembles longer than a Rayleigh length. It is worth considering strategies to phase-match the stationary light over larger distances.

It may be possible to improve phase-matching by coupling to the atoms only over a limited longitudinal region several Rayleigh lengths away from the waist. This would allow the matching of the spherical wavefronts of the control fields to the probe wavefronts. This could work in the special case of roughly equal control and probe wavelengths. Sending copropagating control fields with a much tighter focus would create a uniform control field intensity and therefore coupling off the axis, while also being close to phase-matched away from the waist. However, a tightly focussed control field at the waist of the stationary light is likely to cause problems for whatever interaction is taking place there.

There are two types of completely stationary field which could exist theoretically but not practically in three dimensions. The first is the case of a spinwave with zero momentum, for probe-control pairs of equal wavelength. The spinwave would occupy all space in the transverse directions. Two perfectly mode-matched counterpropagating control fields would generate mode-matched probe fields. Their interaction with the spinwave would then be phase-matched.

The second is in schemes where the probes have higher frequency than the control. A

focussed beam can be composed of probe components all at angles θ_{PM} to the axis with the control fields on axis. For a constant distribution of the components with the axis of propagation as the axis of rotation, this produces a Bessel beam.

8.5.2 Gouy phase and one-dimensional approximations

The Gouy phase can be used as an average for the longitudinal phase mismatch between the Gaussian probe beam and a plane-wave control field. One-dimensional simulations can be run at this phase mismatch to approximate the behaviour of the stationary light. While this would only include the spinwave loss, and not defocussing of the probe and the interplay between the different momenta, it is still a useful method to avoid the computational intensity of higher dimensional simulations.

Alternatively, it is possible to approximately treat the momentum components of a Gaussian mode separately. Again, this is difficult because at larger deviations from the phase-matching angle these will become distributed through momentum space due to the non-stationary behaviour. It may or may not be possible to simply add this evolution up for each component to find the evolution of the entire mode. This may rely on further assumptions about the symmetries of the momentum components and the antisymmetry of the stationary spinwave.

8.6 Conclusion

In this chapter we derived a novel form of stationary light based on a Raman interaction. We showed by a simplified analytic solution how the behaviour of the spinwave under detuned counterpropagating control fields depends on the integral along the length of the spinwave, and how the light is trapped due to an interference between light emitted at different points in the memory.

We then used simulations to show how a more realistic system would behave, including how the various decay mechanisms would affect the stationary light. We also showed that the trapping effect is not perfect for two-dimensional systems (and by extension three-dimensional systems).

Chapter 9

Experimental demonstration of Raman stationary light

Opposition brings concord. Out of discord comes the fairest harmony.

Heraclitus

This chapter presents an experimental confirmation of the Raman stationary light theory. We find an excellent agreement between the behaviour predicted in theoretical models and the observed behaviour.

The published paper relevant to this chapter is

"**Dynamical observations of self-stabilizing stationary light** J. L. Everett, G. T. Campbell, Y.-W. Cho, P. Vernaz-Gris, D. B. Higginbottom, O. Pinel, N. P. Robins, P. K. Lam, B. C. Buchler, Nature Physics 13(1), 68-73 (2016)"

9.1 Experimental Setup

The implementation of Raman SL in rubidium-87 required the use of cold atoms. This was primarily due to the phase matching requirement. While GEM can theoretically be performed on a single Zeeman manifold, it requires good pumping and very clean polarisation. The phase-matching condition for the 6.8 GHz separation of the rubidium-87 ground states resulted in large angles between the control fields and the probe fields. This would lead to prohibitively short lifetimes for the written spinwave in a warm vapour.

9.1.1 Magneto-optical trap for cold rubidium-87 atoms

The theory requires a particular arrangement of laser beam angles and frequencies. In particular, the requirement for the k-vector of the spinwave, $\vec{k}_S$ to be orthogonal to the probe fields. The requirement that the difference of the k-vectors of each pair of probe and control beams had to be roughly identical, $\vec{k}_{p+} - \vec{k}_{c+} \approx \vec{k}_{p-} - \vec{k}_{c-}$ meant some careful alignment techniques had to be considered.

The magneto-optic trap produces a cold, dense cloud of atoms by a position- and velocity-dependent force. A thermal source of atoms is trapped and cooled by the use of magnetic and optical fields. The trap used in the stationary light experiments produces

a long thin cloud with optical depths of the order of hundreds with a temperature of hundreds of milliKelvin. This is the same arrangement used in chapter 4.

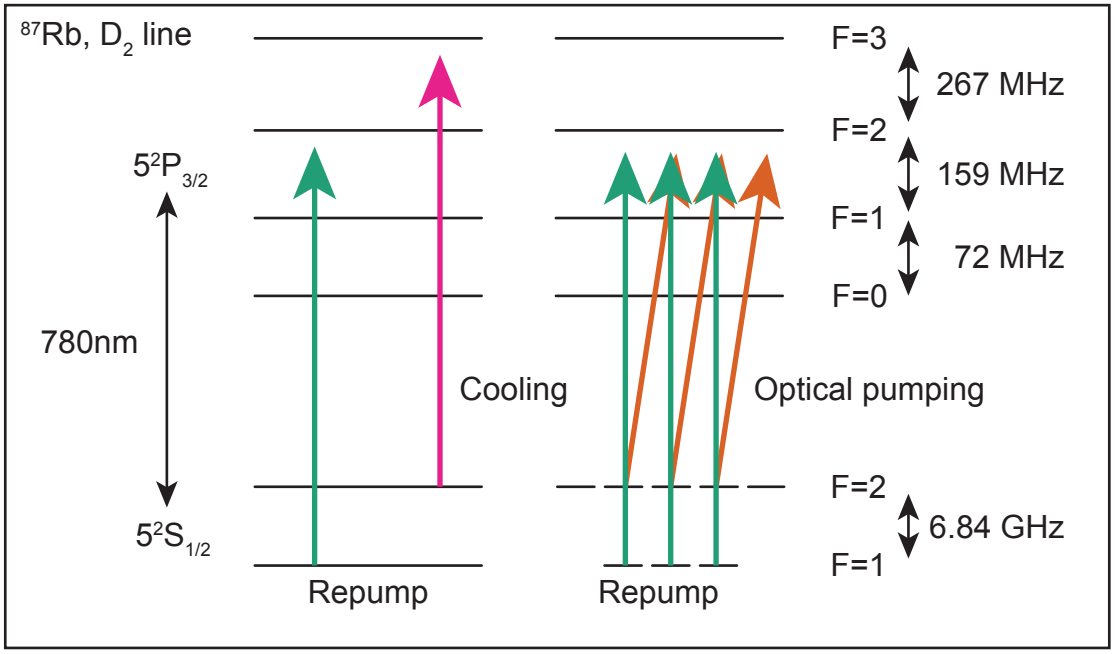

Fig. 9.1. Level scheme used for magneto-optical trap. The cooling beam is detuned from the F=2 to F′=3 transition. The same lasers are tapped off to produce the optical pumping and imaging fields.

9.1.2 Alignment of counter-propagating beam pairs

Aligning all the beams to produce phase-matching and mode-matching was a major additional part of the experiments compared to previous gradient echo memory experiments, so it is helpful to explain our method.

The phase-matching angle θ_{PM} is calculated by finding the angle which equalises the z-components of forward-travelling probe and control momenta.

$$\cos(\theta_{PM}) = \frac{\lambda_c}{\lambda_p} \tag{9.1}$$

$$= \frac{f_p}{f_c} \tag{9.2}$$

$$\approx 1 - \frac{6.835 * 10^9 \mathrm{Hz}}{377.1 * 10^{12} \mathrm{Hz}} \tag{9.3}$$

$$\theta_{PM} \approx 0.00602...\mathrm{rad} \tag{9.4}$$

Fig. 9.2 illustrates the level structure of the rubidium-87 D_1 line, and the transitions addressed in the experiment. The simplified level scheme corresponds to that used in the

theory in chapter 8. The atomic levels correspond to

$$|1\rangle = \left|5^2S_{1/2},\ F=2,\ m_F=+2\right\rangle \tag{9.5}$$
$$|2\rangle = \left|5^2S_{1/2},\ F=1,\ m_F=0\right\rangle \tag{9.6}$$
$$|3\rangle = \left|5^2P_{1/2},\ F=1,\ m_F=+1\right\rangle. \tag{9.7}$$

The effect of the much further detuned level $|3'\rangle = \left|5^2P_{1/2},\ F=2,\ m_F=+1\right\rangle$ is also considered in calculating the transition strengths used in the supporting simulations. In particular, the strength of the Raman absorption is dependent on the optical depth for the transition coupled by the probe field. This is calculated from the measured optical depth of the memory and multiplied by the Clebsch-Gordon coefficients, taken from Steck [111], which account for the transition strengths between particular F and m_f levels. The presence of level $|3'\rangle$ increases the effective optical depth of the Raman transition as the probes have an additional level to couple with.

For incomplete pumping of the atoms to $|1\rangle$, some atoms remain in the level $|1'\rangle = \left|5^2S_{1/2},\ F=2,\ m_F=+1\right\rangle$. The Raman transition $|1'\rangle \rightarrow \left|5^2S_{1/2},\ F=1,\ m_F=+1\right\rangle$ has the same two-photon detuning as the desired Raman transition. That transition can be excited in the case that the control field polarisation is not exactly circular. This population was low enough that it could be ignored in the analysis.

The experiment was set up according to figure 9.3. The probe beam (forward probe) was out-coupled from a fibre and focussed to a waist of about 100 μm at the MOT. The probe was aligned to the MOT by maximising the absorption of the probe near resonance with the transition $|1\rangle \rightarrow |3\rangle$. Imaging of the MOT along the probe direction onto a CCD assisted with this process. After the MOT, the probe was aligned into a detection apparatus. This included polarisation filtering with a polarising beamsplitter and spatial filtering with a pinhole to reduce the amount of control field reaching the detector.

A reference beam with a similar waist was sent in the counter-propagating direction, injected at a non-polarising beamsplitter after the MOT. The reference beam was mode-matched to the probe by maximising the coupling of the reference beam into the fibre carrying the probe beam. The reference beam was aligned into a similar detection apparatus after splitting from the probe with a non-polarising beamsplitter. This was necessary to detect the probe field generated in the backward direction (backward probe) without having to manually align it.

Control fields were aligned based on the probe and reference beams. After the lenses, the control was set a distance from the probe according to the required phase-matching angle. The alignment was fine-tuned using a CCD camera to measure the offsets between co-propagating beam pairs at set locations before and after the MOT.

The alignment was fine-tuned further using the gradient echo memory. Where the k-vectors for the two co-propagating field pairs have some mismatch, the pulses recalled in each direction will be recalled at slightly different times. The z-component of the k-vector affects the rephasing time. This method is only practical where the alignment is near perfect. There are multiple degrees of freedom for the control field modes, and only the one measure of accuracy. An additional, very approximate measure of accuracy was also derived from the backward probe. The angular sensitivity of the backward detection allowed the generated backward probe to be maintained as counterpropagating to the forward probe.

Once the efficiency of the forward and backward detection was accounted for, the control field intensities were balanced using the probe emitted during the stationary light

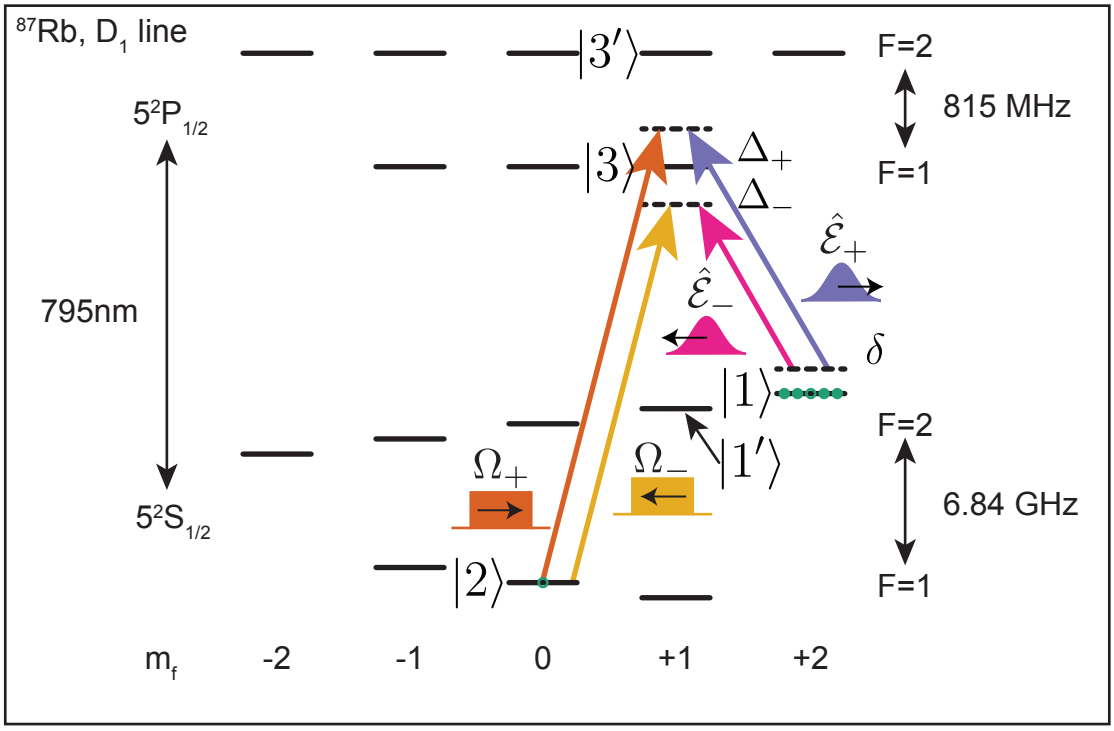

Fig. 9.2. Polarisations and detunings of the experimental optical fields. The atomic energy levels involved in the interaction are shown by $|1\rangle, |2\rangle$, and so on. The forward probe field $\hat{\mathcal{E}}_+$ was stored in the memory with the forward control field Ω_+. In stationary light both control fields $\Omega_\pm$ were turned on, regenerating two counterpropagating probe fields $\hat{\mathcal{E}}_\pm$. Atomic energy levels taken from Steck alkali data [111].

period. The intensities were also fine-tuned by testing whether a spinwave moved in a particular direction during the stationary light period.

9.1.3 Imaging of the spinwave

To assist in understanding the behaviour of the system as well as any experimental limitations, we imaged the spinwave. Observing the spinwave allowed another way to test how well the theory described the behaviour of the system. In principle, it is possible to measure the distribution of the spinwave by recalling it from the memory. In practice, this requires very precise knowledge of all the experimental parameters, and phase sensitive measurements of the optical outputs. It is difficult without this technique to be sure of how well the theory describes the behaviour of the stationary light

The atoms composing the spinwave are in a separate energy level $|2\rangle$. The spinwave is imaged by sending from a transverse direction a large, collimated optical field resonant between $|2\rangle$ and $\left|5P^2_{3/2},\ F=2\right\rangle$, with light taken from the repump beam. This field is absorbed by atoms that have been moved into $|2\rangle$. This gives an incoherent measure of the magnitude of the spinwave.

The resulting illumination was imaged onto a CCD camera (PointGrey Firefly 2.0). Fluorescence from the MOT during the trapping stage was used to ensure the imaging was focussed and the beam illuminated as large as possible a region of the memory.

The magnitude of the spinwave is determined by the absorption of the imaging beam. This is simply done by taking the logarithm of the absorption image and subtracting the logarithm of the reference image. The expected intensity of the imaging beam is $I_0 \exp(-d)$ where I_0 is the intensity measured as the background and d is the optical depth, which is proportional to the number of atoms in the state.

Our imaging system used only one large spherical lens placed at a distance from the MOT to produce approximately 1:3 demagnification. This setup introduced a large barrel distortion in the image. In addition to distorting the spinwave distribution in z and y, this also introduces quantitative errors. The absorption image is compared to a reference image with the same distortion, and regions containing the same density of resonant atoms but with different magnification in the image will appear to have the same density but different size. The same actual spatial distribution of spinwave in one region or another will then appear to have a different number of resonant atoms.

Without a proper reference object for aligning and calibrating the distortion, some error is present in any quantitative measurement. This leaves a significant error margin in the imaging-based measurements of decay, as will be seen in Sec. 9.2.

The imaging data were also used to improve the experimental setup. In particular, a combination of preliminary simulation data using modelled magnetic fields, and imaging data, showed that the magnetic field used to separate the Zeeman levels during the stationary light phase was too inhomogeneous. We were able to increase the size of the magnetic field coils to improve the homogeneity of the bias magnetic field across the atomic ensemble.

9.1.4 Experimental Protocol

Several related experiments were run to acquire optical and imaging data on the behaviour of stationary light. These all involved writing spinwaves similar to those demonstrated in chapter 8. This required that separate areas of the memory be addressed. This

was accomplished using the GEM protocol. By sending two pulses with different carrier frequencies, two spinwaves in different regions of the memory were written. By adjusting the relative phase of the two modulation sidebands, the relative phases of the two regions of spinwave could be changed.

The relative phase of two pulses with different frequency only makes sense when defined at a particular time. In the theoretical example, the spinwave has no phase-winding and each region has a constant phase. This corresponds to a complete rephasing of the spinwave. The corresponding time for the optical pulse is then the input time for the centre of the gaussian envelope. The optical phase of 0 or π is defined by whether the modulated input is maximal or minimal.

The timing and selection of pulses is also important. The stationary light theory requires that we can address separate areas of the memory. This sets some bandwidth constraints on the pulses. The two frequency components should be well separated in the memory so they can be distinguished in the spinwave imaging. The MOT allows for a memory with a time-bandwidth product much greater than 2, so this is possible.

Two pulses with a 20 μs Gaussian time-profile (full width at $I_0 e^{-}2$), each separated in frequency by 70 kHz from the central memory frequency were sent simultaneously into the memory. This was equivalent to sending a single Gaussian pulse, sinusoidally modulated at 140 kHz, which was seen when the pulse was sent directly to the detector.

The pulses were generated as an RF signal by a programmable function generator (Agilent 33620A) and amplified to power a single-pass acousto-optic modulator (AA Opto Electronic MT80-A1-IR).

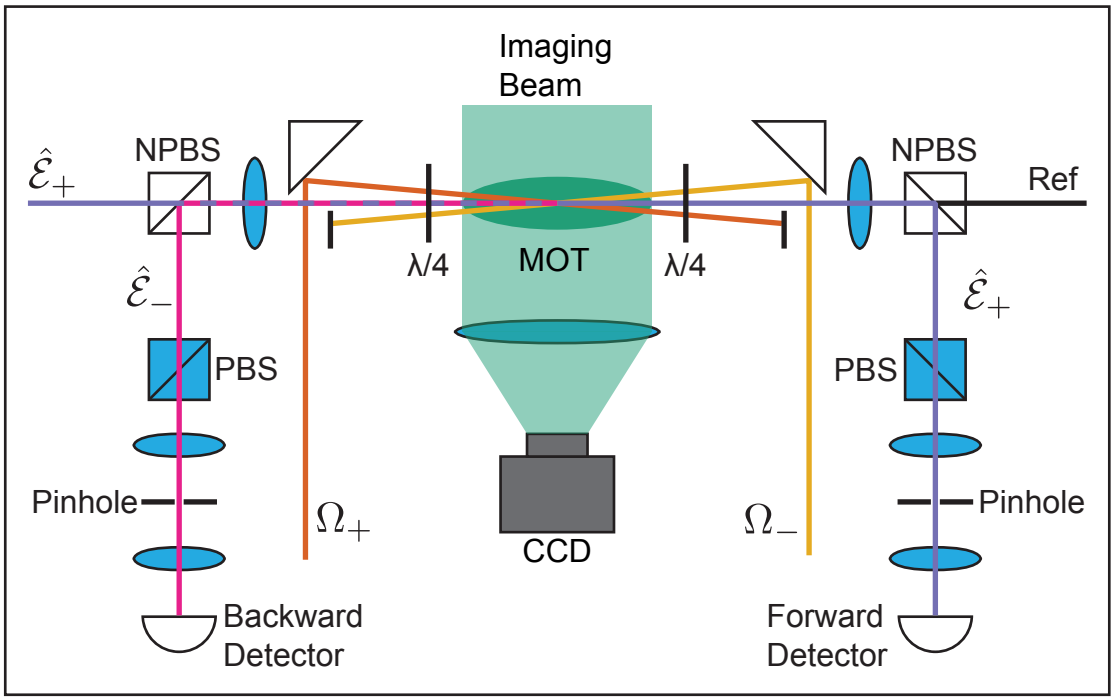

Fig. 9.3. Experimental setup. Beams were combined and split at non-polarising beam-splitters. A large collimated imaging beam illuminated the MOT from the side and was imaged by a single lens onto a CCD.

Once the pulses were stored, the control field was switched off and the gradient was reversed to rephase the spinwave. Then, the gradient was switched off and the counter-

propagating control fields were switched on. For the photodetector data (see the timing diagram in figure 9.4), the counter-propagating fields were left on for 10μs. After another 10μs the gradient was switched back on, and after another 20μs the gradient was flipped and the forward control field switched back on to recall the remainder of the spinwave.

The imaging data (see the timing diagram in figure 9.4) was taken by leaving the counterpropagating control fields on and measuring the evolution of the spinwave. The two pulses were sent at π relative phase, which satisfied the condition that $\int S\mathrm{d}z = 0$; and at 0 relative phase, which did not. A final run of imaging data was taken with a number of different phase increments.

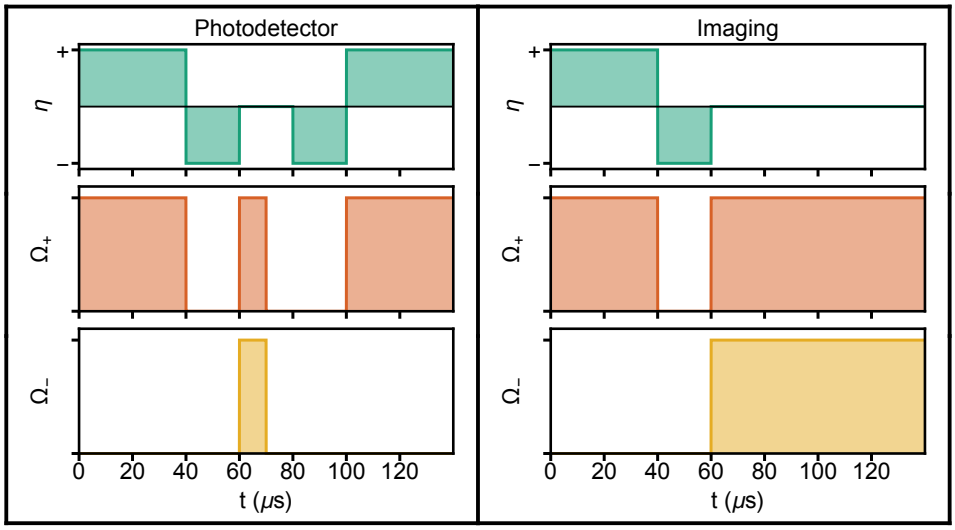

Fig. 9.4. Control field ($\Omega_\pm$) and magnetic field gradient (η) timings for the two stationary light experiments. For the photodetector data, the counter-propagating control fields are applied for 10 μs and the spinwave is then dephased and recalled in the forward direction. For the imaging data, the counter-propagating fields are applied for longer to observe the evolution of the spinwave.

9.2 Results of Raman stationary light experiments

The results are divided into probe input/output data, spinwave imaging, and phase increment imaging data. The optical data and the imaging data confirm the behaviour of the stationary light matches the theory quite well, and the further data is used to analyse some of the experimental imperfections.

9.2.1 Probe input/output data for stationary light

Fig. 9.5 shows the close agreement of the experimental photodetector data with a theoretical prediction from the numerical solution of (7.15). The parameters for the simulations were taken from independent measurements of the atomic ensemble optical depth and the control field intensities and correspond AC-Stark shifts of the Raman lines. The exception is the inhomogeneous decay of the spinwave, too small to measure independently and accurately, which was fitted to the experimental data.

For the stationary spinwave ($\phi = \pi$), there is little light emitted, and the recalled pulse is similar to the recalled pulse where no counter-propagating control fields are applied.

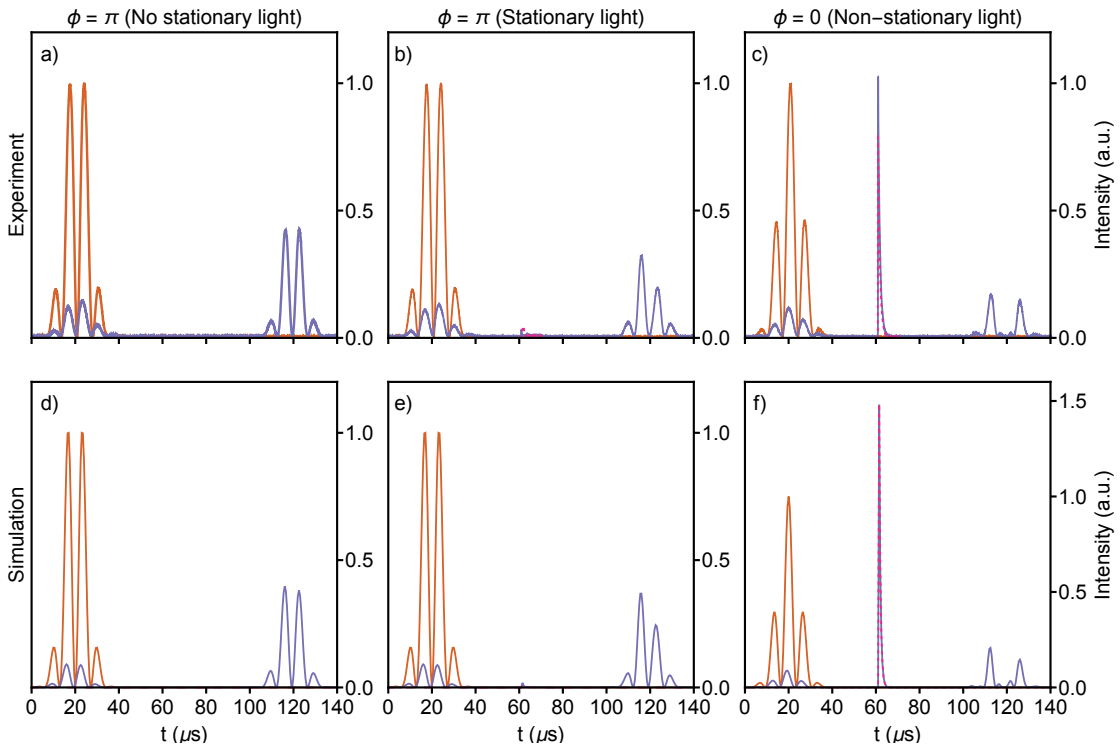

Fig. 9.5. Observation of optical output under stationary light; photodetector data from experiment compared to simulated data. Input (orange), output at forward detector (blue) and backward detector (dashed, purple). The top row displays experimental data and the bottom row shows simulation data.

For the non-stationary spinwave, there is a large spike in the optical output as the spinwave rapidly evolves. The recalled pulse is markedly different from the input reference pulse. The evolution removes the spatially constant component of the spinwave. This component when recalled using the gradient echo memory would correspond to a sinc function in time. The output pulse seen in Fig. 9.5 resembles the input, minus a sinc function which removes most of the central lobe.

The output where no counter-propagating control fields are applied is much more time-symmetric than the output in b). The gradient echo memory is operated under first-in first-out retrieval. Each section of the pulse is stored for an equal time, and experiences roughly equal decay. For b), something happens to the spinwave between the storage and retrieval. Although the two sidebands of the input pulse have equal power, the two regions of the written spinwave have different magnitude due to some incoherent absorption of the probe in the memory. This is difficult to see in the imaging data below, but is clearer in the simulation data. When the two control fields are switched on, the spinwave is redistributed slightly in the memory. This destroys the symmetry of the storage and retrieval.

9.2.2 Imaging results

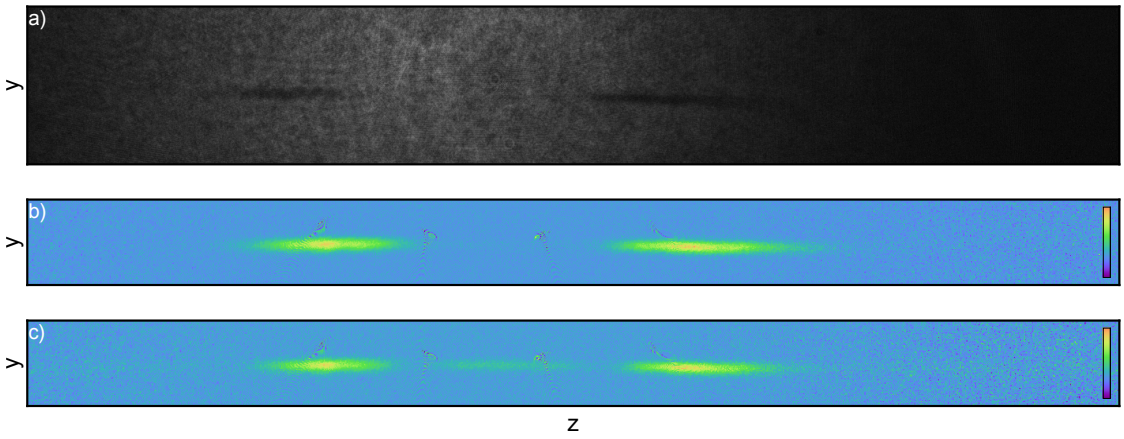

Fig. 9.6. Imaging data. a) raw single camera frame of absorption imaging. b) spinwave magnitude extracted from averaged and distortion corrected absorption imaging. c) processed spinwave after evolution.

In the processed Figs. 9.6 b),c) the colour range is chosen to leave all the data in range. There are a few bright spots in the image that are an artefact of the absorption imaging method and are due to light scattering off corners of the vacuum cell. These take the higher end of the colour range and make it difficult to see the evolution from the images. These could be removed, but instead we improve the visualisation by discarding the unnecessary y dimension. The data as a function of time are presented in Fig. 9.7. By averaging across the y dimension, the SNR is improved further, and the evolution can be more clearly seen as each slice can be placed directly adjacent, rather than separated by empty space. This allows for a comparison between experimental data and simulation data using the same type of plots used earlier in this work to illustrate the spinwave behaviour.

There is some imperfection in the phase-matching or in the intensity matching of the

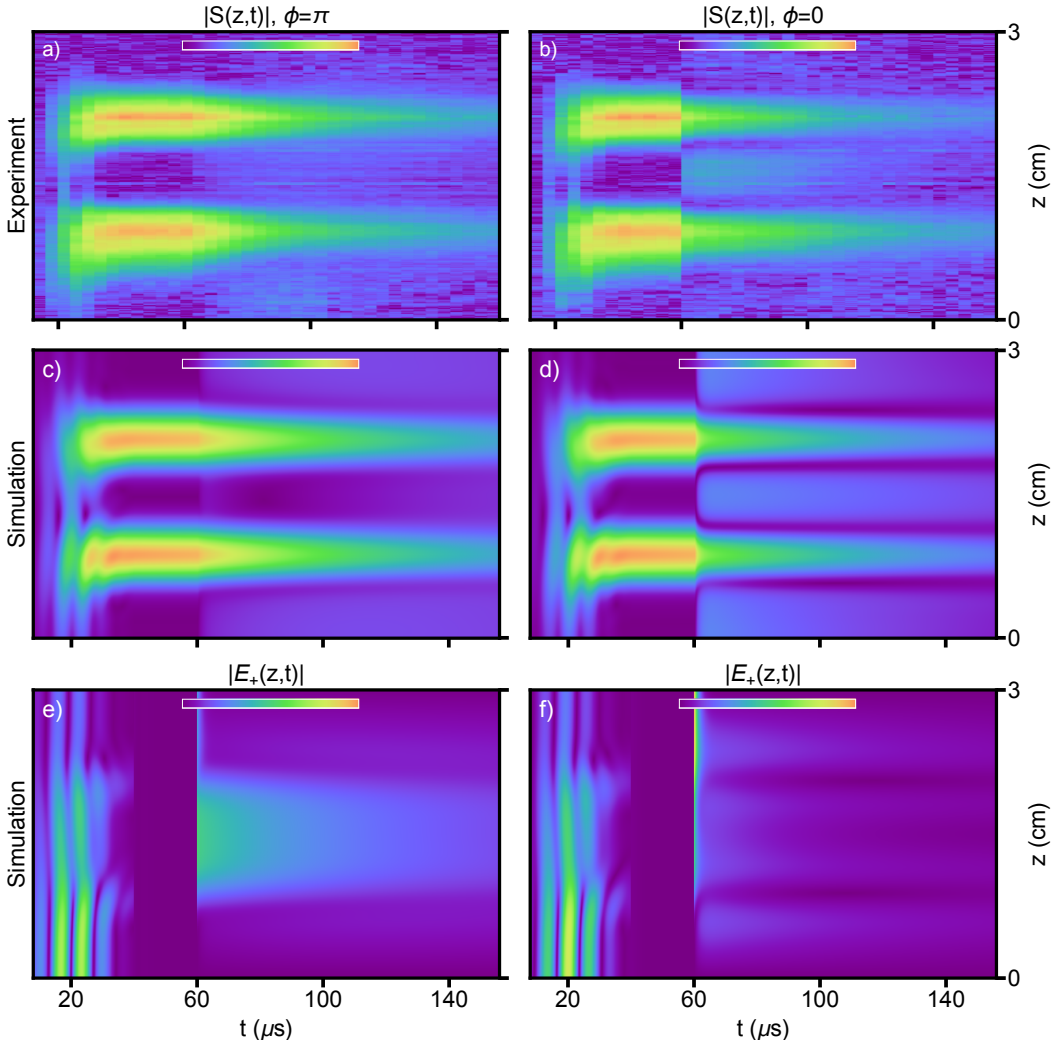

Fig. 9.7. Experimental imaging data and simulated spinwave and probe. a) Stationary and b) non-stationary spinwave magnitude as a function of time. c), d), simulated spinwave magnitude. e), f), simulated forward probe intensities.

control fields. This causes the asymmetry in the evolution of the spinwave. This is most visible in the initially stationary case. This is explored in more detail in Sec. 9.2.4.

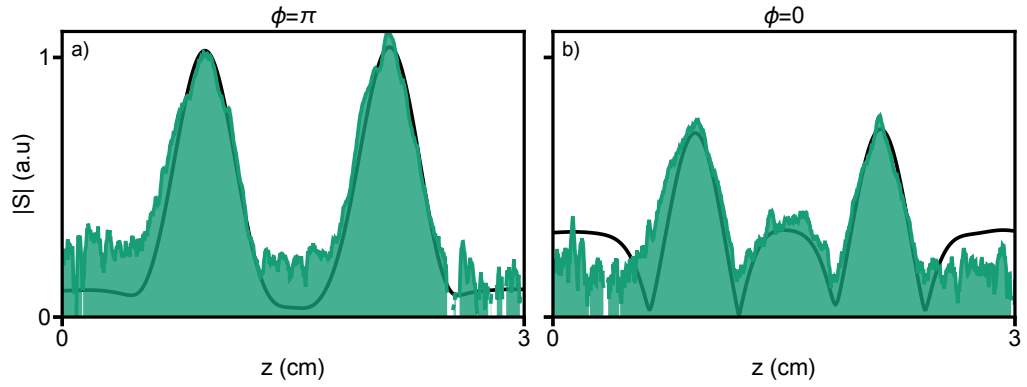

Fig. 9.8. Spinwave after 10 μs of counter-propagating control fields. Initial spinwaves are stationary (a) and non-stationary (b).

9.2.3 Spinwave decay

Imaging is used to measure how well the theoretical decay model agrees with the experimental data. The spinwave is integrated across the whole ensemble and compared to the simulation data, which include the decay terms.

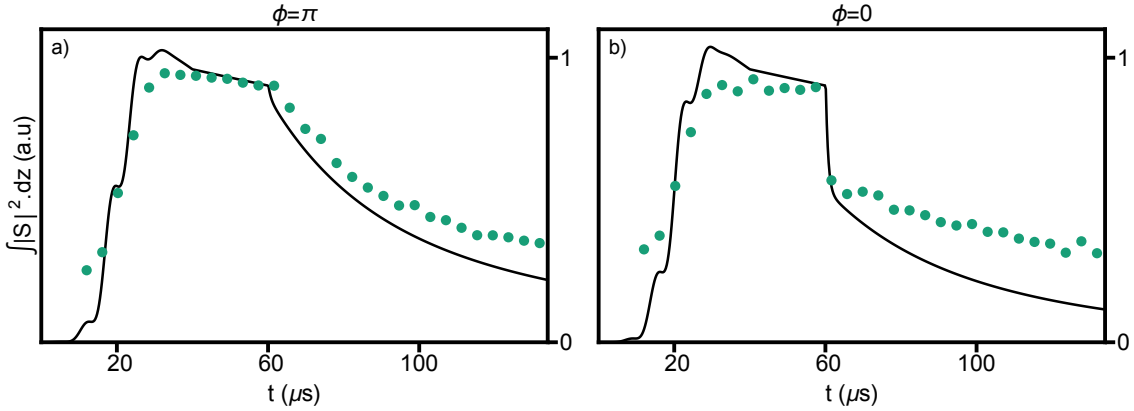

Fig. 9.9. Integrated spinwave amplitude as a function of time for simulation (black line) and experiment (green dots). Spinwaves are initially stationary (a) and non-stationary (b). The scaling between experimental and simulation data is arbitrary but is chosen to emphasise the disagreement between experimental and simulated decay rates after the counter-propagating fields are turned on.

The experimental and simulation imaging data show some differences. In both experimental sets of Fig. 9.9, the spinwave during storage does not peak far above the amount it reaches after storage and before the stationary light. There are several possible explanations. The peak seen in the simulation data at around 30 μs is light that couples into

the spinwave but travels through the memory without being stored. The imaging pulse likely interferes with the storage process by shifting the two-photon resonance. That polariton only temporarily in the memory could be destroyed by the imaging pulse, or the incoming optical field could be prevented from coupling into the memory.

The imaging pulse may also saturate smaller magnitude regions of the spinwave, reducing the absorption and lowering the measured spinwave. This would disproportionately affect the spinwave during times when the spinwave is more spread out, such as during storage and after the counter-propagating fields are turned on.

The qualitative disagreement for the evolution of the non-stationary spinwave is partially due to a different distribution of optical depth between experiment and simulation. In the simulation, the optical depth is assumed to be distributed uniformly along the bandwidth of the memory, whereas in the experiment this is not the case. This may partially account for the disagreement in the behaviour of the spinwaves toward the edge of the memory and the decay rate of the initially non-stationary spinwave.

Considering the range of distorting effects in the imaging data, it is not surprising to see qualitative disagreement with the theory for the non-stationary spinwave.

While the decay mechanisms discussed in the theory section account for experimentally measured decay, and the agreement between theory and experiment is strong, additional undiscovered decay mechanisms for stationary light cannot be ruled out. As will be shown in Chapter 13, the regime necessary to achieve useful nonlinearities with single-photon level states requires optical depths orders of magnitude higher than the current experiment. It is possible that experimental conditions that are optimal for nonlinear interaction will reveal new dynamics.

9.2.4 Phase-matching considerations

Storing pulses with different phases for the modulation gives spinwaves with behaviour in between that of the two cases shown above. This acts like an interferometer; the output during the counter-propagating fields has a sinusoidal dependence on the phase of the input modulation. This dependence can be used to test the behaviour of the stationary light and how well experimental conditions match the theoretical requirements. The data of Fig. 9.7 show a small evolution of the spinwave towards one end of the memory, which can possibly be explained by imperfect phase-matching. Given the difficulty of aligning the control fields it is useful to have additional measures of the quality of the phase-matching.

If the phase-matching is not perfect, the relative phase between probe light produced in each region of spinwave will be different for each direction. This will affect the amount of light seen at each detector, appearing as a phase offset in the sinusoidal dependence for the forward and backward probes. This offset is shown in Fig. 9.10. The minimum and maximum outputs in each direction then occur at different modulation phases of the input pulse, which is detrimental to the stationary light. The total phase mismatch should equal the offset in the phase of the sinusoidal dependence of the backward probe output on the input phase modulation.

There are two explanations for this offset. The first is unmatched dispersion of the probes. The probes are not symmetrically detuned from the second excited level, $|3'\rangle = |5P_{1/2}F = 2\rangle$. According to the Clebsch-Gordon coefficients this is a weaker transition from $|1\rangle$, but it still has some affect at higher optical depths. The dispersion is unbalanced as each probe is red-detuned from this transition.

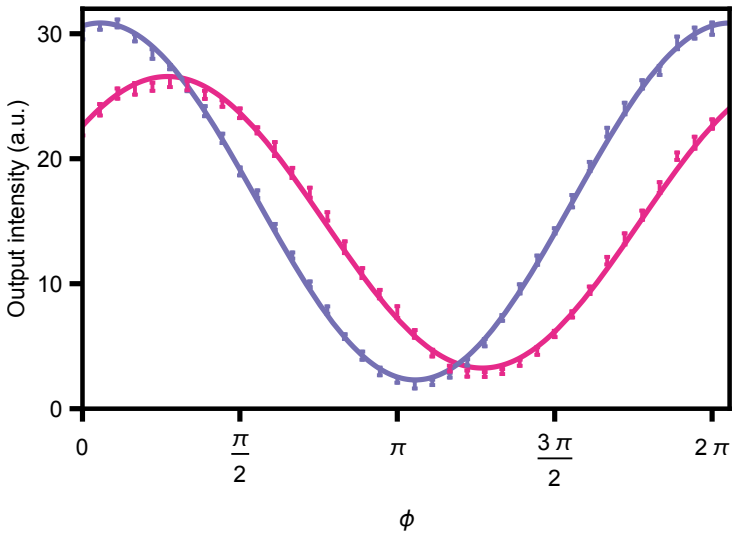

Fig. 9.10. Integrated optical intensity at detector during counter-propagating control fields plotted against input modulation phase ϕ. Forward detector data is in blue, backward in purple, with the solid lines the sinusoidal fits.

The phase offset due to the dispersion is calculated. Assuming a bandwidth of 350 kHz and a total optical depth of 200, the optical depth contained between the peaks, written by pulses with a frequency separation of 140 kHz, is at least 80 for a uniform optical depth, and higher for the experimental optical depth. For the purposes of this calculation an interior optical depth of 120 ± 40 is used. The transition to the second excited state is only $1/3$ as strong, giving an optical depth of about 40 ± 14. The detuning is -815 ± 160MHz for forward and backward probe. The relevant dispersion for the phase-matching is doubled; the backward probe is produced from a spinwave generated by the forward probe, which has experienced the dispersion, and then has to travel again over the region between the spinwaves. The dispersion is then $(0.08 \pm 0.03)\pi$ radian. The phase offset for the backward probe sinusoidal dependence is $(0.16 \pm 0.03)\pi$ radian. The dispersion only accounts for part of the phase mismatch, and can be at least partially compensated by control field alignment.

Another cause of phase mismatch is control field misalignment. The control fields accumulate unequal phase in the z-dimension due to misalignment. The probe fields generated by the different regions of the spinwave no longer have equal phases, producing a direction-dependent offset in the phase of the interference.

In principle, the forward probe output should have no phase offset as it is being recalled with the same control field and under the same dispersive parameters as those under which it was stored. In the experimental data, the forward probe output has a non-zero phase offset due to imperfect rephasing of the spinwave.

Due to the radial symmetry of the probe beams, the phase mismatch due to the use of Gaussian spatial modes cannot be compensated for by control field alignment. We therefore expect this to be the minimum phase offset in the 'interferometer' measurement, unless the sign of other phase mismatch terms allows them to compensate.

The Rayleigh length of the stored probe (waist radius approximately 100 μm) is about 4 cm. The spinwave is written over about 2 cm, amounting to a Gouy phase of about $\pi/4$, which is doubled to give an expected phase offset of $\pi/2$. This is about twice the

measured phase offset. This is a very approximate calculation, and more accurate measurement of the probe waist and spinwave length would allow a quantitative attribution of the various effects to the phase mismatch.

Phase mismatch and strategies to reduce it are treated in more detail in Chapter 11, where the phase mismatch affects the maximum efficiency of a memory based on the Raman-SL scheme.

9.3 Conclusion

We have presented experimental work with Raman-SL. In particular, we have convincingly demonstrated the presence of stationary light in the ensemble, as well as explaining the effects beyond the stationary light equation (8.14). Spinwave imaging was a useful tool developed in the course of this experiment. This allowed diagnosing some magnetic field inhomogeneities and supplementing of the probe input/output data with quantitative measurements of the spinwave evolution. We expect this to be a useful tool in the future, particularly as higher optical depth ensembles allow better SNR and finer resolution images.

Chapter 10

A single-mode atom-optic memory in free space (TRACE memory)

> Mem'ries, may be beautiful and yet,
> What's too painful to remember,
> We simply choose to forget.
>
> Barbra Streisand, *The Way we Were*

Raman stationary light operates by uniform driving of the whole longitudinal mode of an atomic coherence. This interaction can also be used to absorb and re-emit light input to the memory, as the basis of a single-mode atom-optic memory. This chapter sets out the theory behind such a memory. We show how the reduction of the ensemble to a single spatial mode by the control field arrangement allows simple time reversal, and compare the coherent enhancement of the light-matter interaction by the distributed reflection to that of the multiple reflections in a cavity. This gives the name of the scheme, time-reversed and coherently enhanced memory, or TRACE memory.

We find how a memory implemented in this scheme compares to other single-mode memories, and briefly explore the use of a single-mode memory as a temporal mode filter.

10.1 Description of the single-mode memory

The state of a single mode memory under a semi-classical description can typically be described by a single complex number, and is therefore more amenable to a simple theoretical treatment than a continuous, multi-mode memory such as GEM. By making this investigation, the operation of single mode memories will become a little clearer. The impact of the stationary light interaction on the theoretical limits of the memory will also be determined.

To find the best way to operate the memory, we start by examining the recall of a previously stored state. As shown in Chapter 8, the evolution of the spinwave is constant over space. This extends to the situation where the boundary conditions for the probes (that is, the input light) are not zero.

In other words, the memory will emit light if it has already absorbed some light and the control fields are on. By making sure the output probes cancels the input probes, the memory can efficiently store the whole input. From the stationary light interaction, there

are optical fields emitted from both ends of the ensemble, and so two input optical fields are required.

The following equations describe the evolution of a spinwave under the stationary light conditions of Chapter 8, but with the addition of input probe fields from each end of the memory, tuned to two-photon resonance with their respective control fields:

$$\partial_t \hat{S}(\xi,t) = i\sqrt{d}\Gamma\frac{\Omega(t)}{\Delta}\left(\hat{\mathcal{E}}_+(\xi,t)+\hat{\mathcal{E}}_-(\xi,t)\right) \tag{10.1}$$

$$\partial_\xi \hat{\mathcal{E}}_\pm(\xi,t) = \pm i\sqrt{d}\frac{\Omega(t)}{\Delta}\hat{S}(\xi,t) \tag{10.2}$$

where $\hat{S}(\xi,t)$ describes the spinwave and $\hat{\mathcal{E}}_\pm(\xi,t)$ describe the optical fields travelling in the increasing (+) and decreasing (-) ξ directions, with the memory parameters of optical depth d, excited state half-linewidth Γ, one-photon detuning Δ, and control field Rabi frequency Ω.

As in Eq. (8.14), the evolution of the spinwave is dependent on the integral of the spinwave, but here there are also the non-zero input fields $\hat{\mathcal{E}}_+(z=0,t)$ and $\hat{\mathcal{E}}_-(z=1,t)$:

$$\begin{aligned}\partial_t \hat{S}(t) = &-d\Gamma\frac{\Omega(t)}{\Delta^2}\int_0^1 \hat{S}(\xi',t)\mathrm{d}\xi' \\ &+ i\sqrt{d}\Gamma\frac{\Omega(t)}{\Delta}\left(\hat{\mathcal{E}}_+(\xi=0,t)+\hat{\mathcal{E}}_-(\xi=1,t)\right).\end{aligned} \tag{10.3}$$

Equation (10.3) is obtained by inserting Eqs. (10.2) into (10.1). The input probe fields are the boundary conditions for the probe field integrals. Where the input fields are zero and a constant control field is applied, an exponential decay of the spinwave occurs simultaneously with equal probe fields exiting both sides of the memory.

The absence of a spatial degree of freedom simplifies the analysis of the memory based on a time-reversal symmetry. The correct output coupling profile to achieve the desired pulse can be found by time-reversing the coupling profile used to store that pulse. We set $\hat{\mathcal{S}}$ constant in z, assuming that any spinwave has been written with the counter-propagating control fields and is constant in space.

The time-reversal symmetry can be fully exploited when $\hat{\mathcal{E}}_+(0,t) = \hat{\mathcal{E}}_-(1,t) = \hat{\mathcal{E}}_{IN}$.

By sending the input $\hat{\mathcal{E}}_+(0,t) = \hat{\mathcal{E}}_-(1,t) = -i\sqrt{d}\Gamma\frac{\Omega(t)}{\Delta}\int_0^L \hat{S}(z')\mathrm{d}z'$, the output fields $\hat{\mathcal{E}}_+(1,t) = \hat{\mathcal{E}}_-(0,t) = 0$ and $\partial_t\hat{S} = \frac{\Omega(t)}{\Delta}\int_0^1 \hat{S}(\xi')\mathrm{d}\xi'$, the negative of the original. For efficient storage, the time evolution of the input fields must match the time evolution of the spinwave. Therefore, an exponentially rising pulse of the profile $\hat{\mathcal{E}(\sqcup)} \propto \exp\left(d\Gamma\frac{\Omega^2}{\Delta^2}t\right)$ is efficiently absorbed. When the pulse is complete, the coupling fields may be turned off to prevent re-emission, and turned on at a later time. By adjusting the time profile of the control fields, other pulse shapes may be absorbed efficiently.

For $\hat{\mathcal{E}}_{IN} = -\hat{\mathcal{E}}_{OUT}$ the optical fields are completely absorbed, as in Fig. 10.1 a). The optical field $\hat{\mathcal{E}}_{OUT}$ depends on the previous optical field sent, as well as the coupling parameters. By assuming all the previous input field was efficiently absorbed, an expression

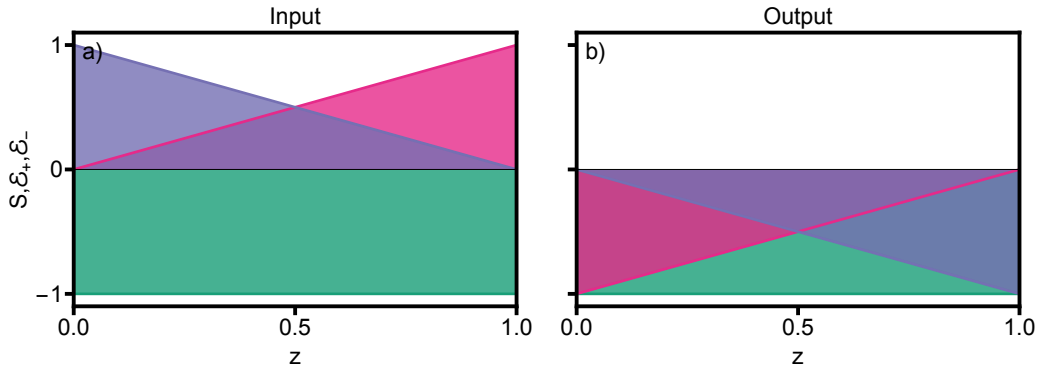

Fig. 10.1. Time reversal in the TRACE memory. a) The forward input probe field $\mathcal{E}_+$ (blue) is absorbed as it travels through the memory, as the probe field generated in the memory cancels it out by the end. Similarly for the backward probe $\mathcal{E}_-$ (purple). The sum of the fields is equal throughout the memory so there is a spatially constant evolution of the spinwave. b) The input is set to zero, and the fields emitted in the memory are not cancelled out. These fields leave the memory while causing an opposite evolution of the spinwave, effectively time-reversing the storage process.

for the output field can be found:

$$\hat{\mathcal{E}}_{OUT}(t) = i\sqrt{d}\frac{\Omega}{\Delta}\int_0^1 \hat{S}\mathrm{d}z' \tag{10.4}$$

$$= -d\Gamma\frac{\Omega^2(t)}{\Delta^2}\int_0^t |\hat{\mathcal{E}}_{IN}(t')|\mathrm{d}t' \tag{10.5}$$

The absolute value of $\hat{\mathcal{E}}_{IN}$ appears due to an arbitrary phase established at the start of the pulse. The output at the very start of the pulse will be of opposite phase to the input. There is no absolute phase requirement for the control field but the control field phase must be changed as the phase of the input changes to maintain the cancellation of input and output.

Then, we can solve for the desired coupling parameter. $\Omega(t)$ is the obvious choice:

$$\Omega(t) = \frac{\Delta}{\sqrt{d\Gamma}}\frac{\hat{\mathcal{E}}_{IN}(t)}{\sqrt{2\int_0^t |\hat{\mathcal{E}}_{IN}(t')|^2\mathrm{d}t'}} \tag{10.6}$$

This equation places limits on the time profile for $\hat{\mathcal{E}}_{IN}(t)$: It must not result in $\Omega(t) \to \infty$ or $\partial_t\Omega(t) \to \infty$ as these would not be experimentally feasible to generate. Given that the control field is proportional to the probe field divided by the time integral of the probe field, the pulse bandwidth is restricted by how much control field power is available. A more fundamental limitation is that at very large coupling field strengths the two-level approximation will break down.

Some examples of well-behaved inputs are $\hat{\mathcal{E}}_{IN}(t) = e^{kt}$, giving constant $\Omega(t) = k\Delta/\sqrt{2d\Gamma}$, and Gaussian inputs which have a more complicated control field profile. Gaussian inputs are also studied in more detail in Sec. 10.3.

Calculations including loss will be similar to those for other single mode memories [93], so we will not include them here. A significant difference is that in the ideal three-

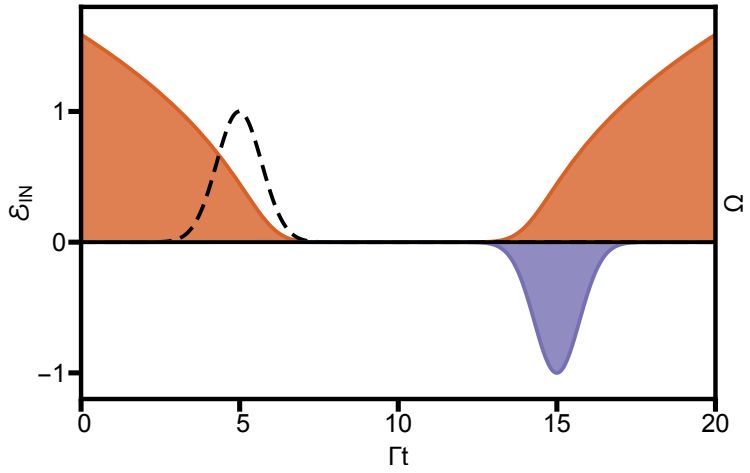

Fig. 10.2. Storage and retrieval of a Gaussian pulse with TRACE. The input pulse ($\hat{\mathcal{E}}_{\mathrm{IN}}$, dashed) and output pulse ($\hat{\mathcal{E}}_{\mathrm{OUT}}$, blue) are flipped in phase where the coupling field (Ω, orange) is not.

level case the AC-Stark shift will not modify the phase (so the phase of the control fields will not need to be continually modified depending on their intensity).

10.1.1 Free space

Single-mode memories in ensembles operate by ensuring a uniform interaction along the length of the memory. As discussed in Chapter 3, this is typically done by placing the ensemble in a cavity. The input field traverses the ensemble many times as it is absorbed, and so is absorbed equally along the length of the ensemble.

In TRACE, the input fields are not absorbed equally along the length of the memory. Instead, each field is absorbed entirely during one traversal of the memory, and the sum of the driving of the spinwave due to the two input fields produces the spatially uniform spinwave.

This avoids some of the complications of using a cavity. The input field needs to be resonant with the cavity, so the mirrors need to provide an optical path length with wavelength level stability. A back action on the memory due to light circulating in the cavity also needs to be accounted for. That is, light that has been emitted from the ensemble can be reabsorbed while it is trapped in the cavity. This is not an issue in TRACE as the light does not travel through the memory again once it is emitted.

Our scheme requires a good mode overlap between the two counter-propagating input probe fields and the re-emitted probe fields, with the maximum efficiency being multiplied by this overlap. It also requires phase stability between the co-propagating pairs of probe and control fields. This is similar to the requirement for a stable optical path length in a cavity, but allows for different methods of satisfying the requirement.

The fact the interaction is entirely controlled by optical fields that are generated far from the ensemble may also allow more flexibility with the memory. For example, multiple spatial modes might be accessed in the same ensemble by multiple sets of control fields, which would be more difficult if a cavity were required for each spatial mode.

The stationary light equations (8.14) are based on the assumption that no probe light

is input during the stationary light period. Then, all light generated in the memory has a relative phase dependent on the phase-matching of the control field and the relative phase encoded in the spinwave. This condition breaks down when light is input during the stationary light period. There is no expectation that the phase of light entering the memory matches the phase of light being generated within the memory, or that the phase of the input from each side (and the corresponding spinwave) will match.

Looking back to the stationary light case, we saw that when there was a π phase between the counter-propagating probe fields, the spinwave did not evolve. The same applies here, but this time it is an unwanted effect. If no spinwave is generated, the probe fields are not absorbed and the storage fails. Therefore it is necessary to control the relative phase of the input fields to run an efficient memory. Providing there is a sufficiently stable optical path, this can be accomplished by adjusting the phase of one of the control fields.

10.2 Efficiency of TRACE memory

The scattering of the control field off atoms in state $|2\rangle$ is an unavoidable source of loss in three-level memories. We can show how this sets a fundamental limit to the efficiency of the TRACE memory. The third level must be included in the equations to quantify the loss. A complication is that the excitation of the ensemble into the third level is separated into two coherences, each one interacting with only the forward or backward travelling probe. To maintain these as constant in space, the adiabatic approximation must be strictly satisfied.

To simplify equations, we constrain $P_+ = P_-$, and P must be constant in z. For a zero input, $\mathcal{E}_{IN\pm} = 0$,

$$\begin{aligned}
\mathcal{E}_{OUT\pm} &= \mathcal{E}_{IN\pm} + i\sqrt{d}P_\pm \\
\partial_t P_+ &= -(\Gamma + i\Delta)P_+ + i\Omega S + i\sqrt{d}\Gamma\mathcal{E}_{OUT} \\
&= -\left(\Gamma(1+dz) + i\Delta\right)P_+ + i\Omega S \\
\partial_t P_- &= -\left(\Gamma(1+d(1-z)) + i\Delta\right)P_- + i\Omega S \\
\partial_t S &= i\Omega^*(P_+ + P_-)
\end{aligned} \tag{10.7}$$

Then the total time derivative can be found.

$$d/dt(|S|^2 + |P_+|^2 + |P_-|^2) = -2\Gamma(|P_+|^2 + |P_-|^2) + 2d\Gamma(z|P_+|^2 + (1-z)|P_-|^2) \tag{10.8}$$

In the adiabatic limit, $P_+ = P_-$. Substituting in the output probe fields for the coherences $P_\pm$, and integrating gives

$$\int_0^\infty |\mathcal{E}_{OUT}|^2 \mathrm{d}t = \frac{d}{\Gamma(d+2)}\left[|S(0)|^2 + |P_+(0)|^2 + |P_-(0)|^2 - |S(\infty)|^2 - |P_+(\infty)|^2 - |P_-(\infty)|^2\right] \tag{10.9}$$

For some control field profile, the remaining coherence at $t = \infty$ is zero. The maximum retrieval efficiency is then $d/(d+2)$. For comparison, this gives an error of approximately $2/d$ for large d. For comparison, the error for a cavity is $1/C$, where the cooperativity C is the effective optical depth equal to the average number of times light

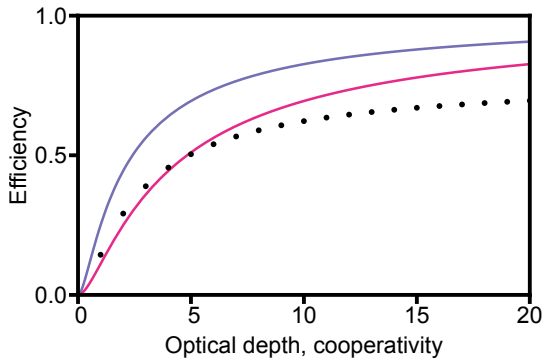

Fig. 10.3. Efficiency for cavity single-mode memory as a function of cooperativity (blue line), efficiency for TRACE as a function of OD (purple line) and approximate optimal efficiency of Raman memory based on simulations (black dots).

traverses the cavity multiplied by d, and for a single directional input free-space memory the error approaches $2.9/d$ at large d. TRACE can therefore outperform a single directional free-space memory, while requiring a more complicated arrangement of the optical fields. A comparison of optimal efficiencies depending on optical depth or cooperativity is shown in Fig. 10.3. At d or C equal to 20 these efficiencies are already very close to their asymptotic limits, while there is a crossover point at $d = 5$ where TRACE starts to outperform a Raman memory.

The requirement that $P_+ = P_-$ places restrictions on bandwidth to a similar degree as for a cavity memory. However, there is no possibility of directly implementing 'fast' memory as described in [93], where no control field is applied during the input pulse and the input is a rising exponential $\exp(d\Gamma t)$. The input is absorbed entirely into the coherences $P_\pm$, which become non-uniform in space under the non-adiabatic condition. This non-uniformity results as the optical coherences are each driven entirely by a single probe field, rather than the probe field plus the recall of the spinwave S. The non-uniformity in $P_\pm$ results in an inefficient out-coupling, reducing the efficiency of the memory.

This problem can be avoided in ordered media. Each field can couple to the same optical coherence if the periodicity of the medium matches the wavelength of the input field. This is because the excited emitters simply cannot store information to distinguish between the direction of the fields they absorb. In fact, this situation allows for the use of a single control field, increasing the efficiency of the memory and removing the above bandwidth limitations.

10.3 Temporal mode selection in single-mode memories

The storage of an optical pulse in a single-mode memory depends on a destructive interference of the input and output fields. This means that only temporal modes identical to the desired storage mode will be stored efficiently, and some input modes may not be stored at all. We are curious to test the limits of just how selective a single-mode memory can be. We use as an example a set of input pulses with Hermite-Gaussian (HG) time profiles.

We ignore losses in order to simplify the equations. We call the input signal field $\mathcal{E}_{IN}(t)$, the mode of the memory $S(t)$ and the coupling field $\Omega'(t)$, which determines the coupling strength into and out of the memory. We then have the output $\mathcal{E}_{OUT}(t) =$

$i\Omega' S(t)$, and

$$\partial_t S = i\Omega'(\mathcal{E}_{IN} + \mathcal{E}_{OUT}) \tag{10.10}$$
$$\mathcal{E}_{OUT} = \mathcal{E}_{IN} + i\Omega' S \tag{10.11}$$

Efficient storage corresponds to setting $\mathcal{E}_{OUT} = 0$. These equations can be used to find the time profile of $\Omega(t)$ that efficiently stores a particular $\mathcal{E}_{IN}(t)$, with the result ([93]) that

$$\Omega'(t) = \frac{\mathcal{E}_{IN}(t)}{\sqrt{2\int_0^t |\mathcal{E}_{IN}(t')|^2 \mathrm{d}t'}} \tag{10.12}$$

These equations ignore loss due to control field scattering and other decay mechanisms, and so do not give realistic efficiencies. However, they are useful for finding the approximate control field profiles suited to a particular input.

10.3.1 Hermite-Gauss polynomials

The HG modes are a set of orthogonal modes that can be used as a basis set to describe the time profile of any memory input or output. We have seen already that an input that is Gaussian in time can be stored using a well-behaved control field time profile. This is also true for the Hermite-Gaussian polynomials in general. We use the HG modes as an example of single-mode memory selectivity by simulating the storage of each of the first ten modes with each of the control field time profiles optimised for these modes.

We have seen in Fig. 10.2 that Eq. (10.12) has a smooth form for a Gaussian input $\mathcal{E}_{IN}(t) = \exp\left(-(t-t_1)^2\right)$, giving the control parameter

$$\Omega'(t) = \frac{\exp\left(-(t-t_1)^2\right)}{\sqrt{\mathrm{Erfc}(\sqrt{2} * (t-t_1)) * \sqrt{\pi/2}}}. \tag{10.13}$$

To retrieve the pulse in the same temporal mode, we simply reverse the coupling field time-profile (In Gorshkov et al. [93], this is how Eq. (10.12) is derived). Putting these equations for the input and coupling fields, with $t_1 = 5$ and retrieval time $t_2 = 15$ in a simulation gives the results in Fig. 10.2.

We try also multiplying a Gaussian input by the Hermite polynomials (for $t - t_1 \equiv x$):

$$\begin{aligned} HG_0 &= 1 \\ HG_1 &= 2x \\ HG_2 &= 4x^2 - 2 \\ HG_3 &= 8x^3 - 12x \\ HG_4 &= 16x^4 - 48x^2 + 12 \end{aligned} \tag{10.14}$$

Efficient storage of the higher-order polynomials with the appropriate time profiles for the control fields operates as expected. Fig. 10.4 shows two-level simulation results for HG_0 and H_1. The control field time profiles remain smooth and low-valued. Storage of HG_0 with HG_1 protocol leads to the emission of a time-shifted pulse, with zero storage efficiency at the end of the input control field.

We can see also see in Fig. 10.4 that storage of HG_1 with the HG_0 control profile transforms it to an approximately Gaussian or HG_0 mode, slightly time-shifted from the

central input time of $t = 5\mu s$. This is very closely matched to the H_0 mode with an overlap of $\langle HG_{1OUT}|HG_0\rangle > 0.997$ according to simulated data.

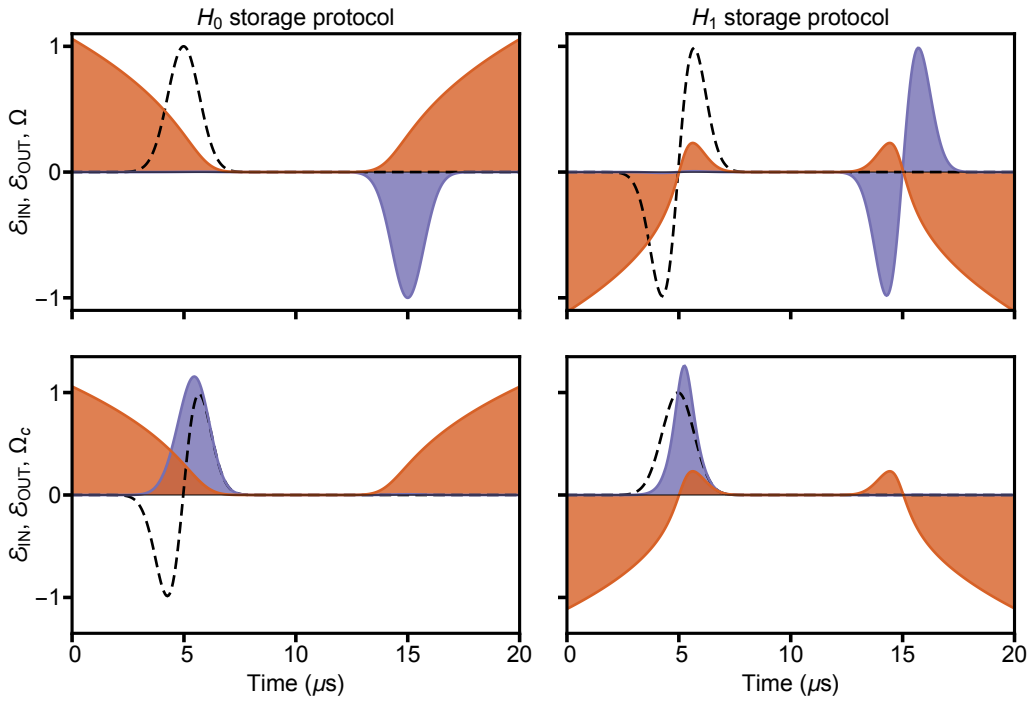

Fig. 10.4. Simulated storage and retrieval of different pulse profiles with TRACE. Input pulses (dashed) and output pulse (blue) and coupling fields (orange). Each mode is stored efficiently under its own storage profile and not stored at all under the other control field profile.

Fig. 10.5 shows the efficiencies for the storage of the first ten HG modes with each of the different optimal coupling profiles for those modes. Each mode is efficiently stored under the coupling profile for that mode. Odd HG modes are not stored at all under even coupling profiles and vice versa, giving the plot its checkerboard pattern. The unstored modes are not converted to other HG modes, except in the special case of HG_1.

The efficiency of storing input x under control field timing y is always within 10^{-6} of the efficiency of storing mode y under control x, at least for the first ten modes. This gives Fig. 10.5 its transpose symmetry. This is unlikely to be a coincidence but we are unable to offer any compelling reason for it to be the case.

Another unexplained feature of Fig. 10.5 is that the efficiencies are all rational numbers. For example, the efficiencies in the first non-zero diagonal after the central diagonal are $\frac{1}{3}, \frac{3}{5}, \frac{5}{7}, \frac{7}{9}$, and so on.

10.3.2 Limits to mode selectivity

For the set of HG modes, the first two modes are not stored at all under each others control field profile. One is even about the centre of the pulse and one is odd, a pattern predicting no storage that is obvious in Fig. 10.5. This cannot produce orthogonality as the number of modes is increased beyond two. We can also rule out obtaining three orthogonal modes by stretching or squashing the envelope of HG_2 in time so that it is orthogonal with HG_0. We do not expect that it is possible to find a set of modes larger

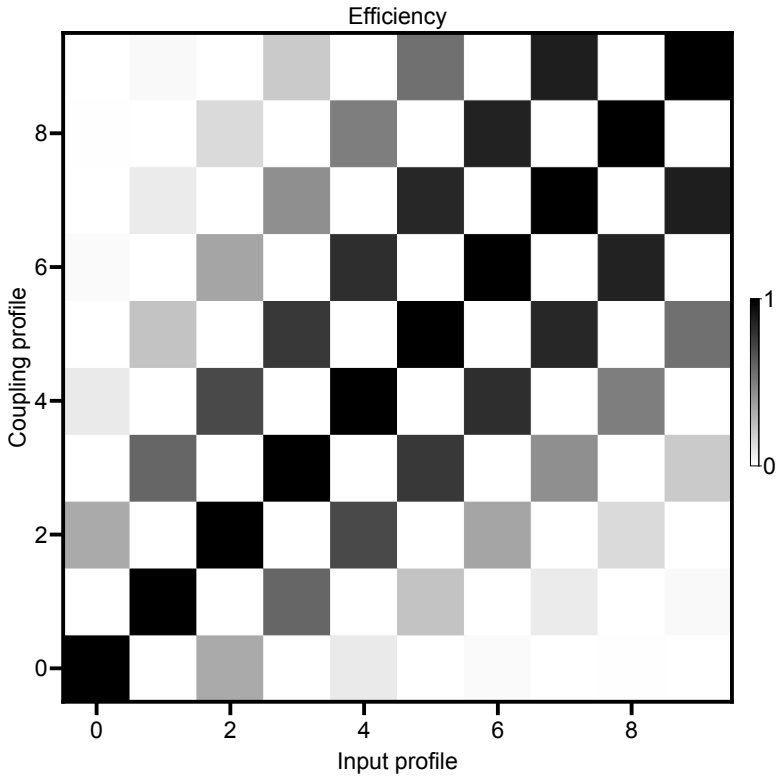

Fig. 10.5. Efficiency of storage of input profiles corresponding to HG_0-HG_9 temporal modes under the optimal control field timings for the same set of modes. The central diagonal corresponds to modes stored under their corresponding control field profile and have the highest efficiency.

than two that is orthogonal with respect to storage.

Another set of modes that can be considered in terms of orthogonality are rising exponential inputs. By separating the inputs in time, and using the time the coupling fields are on to select different inputs to store, some amount of selectivity is possible. This amount will depend on how far the inputs are separated in time as the beginning of the rising input will always be stored efficiently if the coupling fields are on. As the inputs are moved closer together in time, each mode will be stored more efficiently by the other coupling profiles.

This effect can be seen in Fig. 10.5, even though that is for a different set of modes. The higher HG modes are stored more efficiently for nearby coupling profiles compared to the lower modes. This is because the lobes of the HG functions are moved closer together, resulting in a greater overlap with the coupling profile.

10.4 Conclusion

We have shown how the evolution of the spinwave in Raman stationary light allows for the implementation of a temporally single-mode memory. This memory has a comparable efficiency to a single-mode memory implemented in a cavity, without the extra bandwidth limitations. We also investigated the mode selectivity of a general single-mode memory, showing that at least a set of two modes sent at one time can be made orthogonal.

Chapter 11

TRACE: Experimental demonstration

This chapter presents data from an experimental implementation of the scheme of Chapter 10. We demonstrate a proof of concept of the TRACE memory. The experimental efficiency is limited by alignment and phase-matching issues. We use these outcomes to propose a more robust setup.

11.1 Experimental setup

This experiment used a nearly identical setup and laser frequencies to the stationary light experiment in Chapter 9, with the following modifications.

1. The optical depth of the MOT at the probe transition was significantly improved to up to 500 ± 100 by using a cooling and compression sequence optimised with machine learning ([156]).

2. The one-photon detunings $\Delta_{\pm}$ were increased to ± 230 MHz to reduce the incoherent absorption. The theoretical probe transmission absent control fields is $\mathcal{E}_{+}(z=L) \approx 0.9998\mathcal{E}_{+}(z=0)$

3. The reference beam was replaced with an *input* probe $\mathcal{E}_{-}$ tuned to $-\Delta$ from the transition $|1\rangle \rightarrow |3\rangle$.

4. The mode-matching was improved by placing identical input optics including fibre couplers and mode-matching lenses in the backward and forward probe paths.

5. The phase-matching (the elimination of dependence of relative phase between the spinwave couplings on longitudinal position) was improved using the alignment techniques detailed below.

All the beams had separate optical paths, and the control fields were combined with their probe fields on wedge mirrors. This meant that the relative phase fluctuated and was essentially random between each cycle of MOT compression, but was more stable over 100μs.

11.1.1 Phase-matching techniques

The experiment was initially aligned using similar methods to the stationary light experiment. The probe beams were mode-matched and the control field angles were set using

the displacement from the probe beams on CCD cameras placed along the beam path at known distances from the MOT. All fields were sent into the MOT at the same time, for a period of 100 μs. By using a significantly larger probe power than that used while taking memory efficiency data, visible images were produced with the CCDs. These images were from the probe field transmitted through the memory as well as the counter-propagating probe field reflected from the memory. The overlap of these two images corresponded to a good starting point for alignment of the phase-matching of the fields.

A second technique involved adjusting the control field alignments while running the experiment. Each control field was aligned both in incident angle and displacement in both transverse dimensions to find the optimal efficiency. This process was iterated while adjusting the two-photon detuning and control field intensities to optimise the efficiency, as moving the beams changed the control Rabi frequency at the ensemble.

A third technique was inspired by the phase-matching measurement described in Sec. 9.2.4. The relative phase between the two input pulses was cycled through over a single experimental run (i.e. after a single MOT compression cycle) by sending 30 input pulses with a 2 KHz offset in the two-photon detunings. This produced four sinusoidal signals, one for each input and output pulse in each direction. The phase offsets between these signals corresponded to the accuracy of the phase-matching. A zero offset between each pair of input (output) signals and a π offset between input and output corresponded to the best phase-matching. This measurement was optimised by adjusting one or both control field alignments. This technique required reasonable efficiency, and so was dependent on having already aligned the control fields and optimised intensities according to the previous techniques.

As with the stationary light experiment, the phase-matching was difficult to reproduce from day to day. The dependence of the phase-matching and reflection efficiency on a large number of degrees of freedom made finding the optimal alignment difficult. A more accurate phase-matching was required for this experiment compared to the stationary light experiment, as the ensemble length over which the phase needed to be matched was increased and the phase-matching also involved mode-matching the two input fields. This was the largest factor in making this experiment a proof of concept for the TRACE memory, rather than a demonstration of a higher efficiency memory at low optical depths compared to a multi-mode memory.

11.1.2 Detector calibration

The measurement of two output fields with two separate detectors required more careful calibration than the stationary light experiments. The efficient operation of the memory required equal intensities of the two input fields. The reliable measurement of the memory efficiency also required an accurate measurement of the relationship between the optical power at the detector or at the memory output (due to the different losses in the two optical paths) and the voltage output of the detector.

The different losses were accounted for by measuring more powerful probe beams directly after the MOT chamber and before the detector. The different sensitivities of the detectors were then accounted for by sending identical optical powers incident on each detector. This produced a single number that could be used to scale the pulses measured at the detectors, ensuring that identical pulse energies were sent in each direction.

11.1.3 Storage protocols

In this experiment, we stored probe pulses with exponentially rising envelopes. This input was chosen so that control fields with constant intensity could be used, simplifying the control of the experiment. Time-dependent control envelopes complicate the optimisation of the control intensity. Also, because of the presence of multiple excited states, the two-photon detuning was still dependent on control field intensity (the AC-Stark shifts do not cancel in the experiment as they do in the three-level theory). Using time-dependent control envelopes would require the use of time-dependent phase and frequency profiles to ensure the relative phase of the spinwave couplings remained constant over the duration of the input pulse.

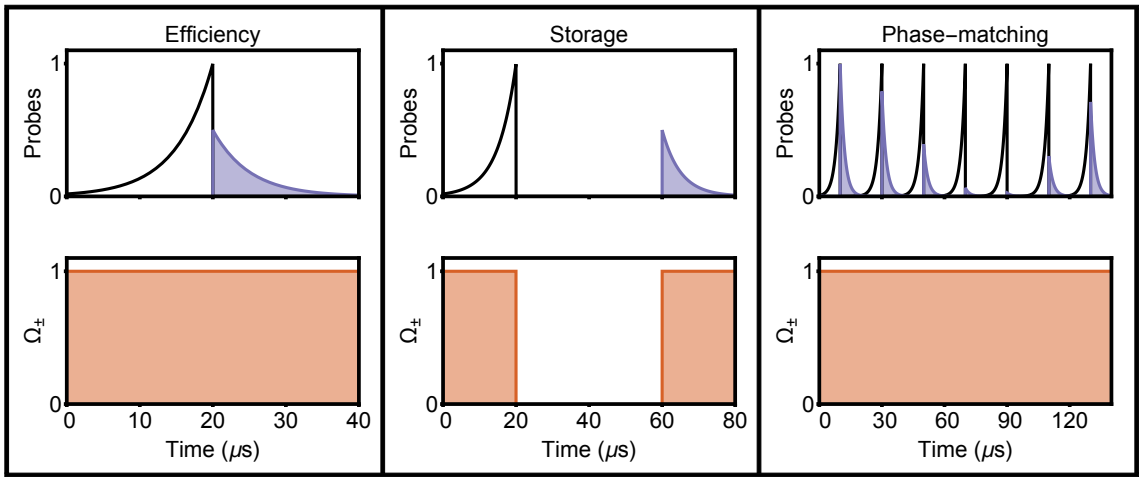

Fig. 11.1. Input/output pulse and control field timings for the experimental datasets. Input probe (black lines) and representative output pulses (filled blue lines), along with control field timings (filled orange lines) for the different experimental protocols.

The memory was characterised with three protocols:

1. A single pair of input pulses were sent over 2000 runs to gather statistics on overall efficiency and mode-matching of the inputs. The relative phase between the two inputs was random with each shot, so some input pulses destructively interfered and were not stored. However, all the shots contributed to the statistics, as the amount of destructive interference depended on the mode-matching and intensity matching of the inputs.

2. The control fields were turned off immediately after the inputs were sent and turned on again some time later in order to store the pulses for a set time. The characteristic decay time of the memory was measured.

3. 17 pulses were sent in each experimental run, over which the relative phase appeared relatively stable. The phase between the spinwave couplings was incremented with each input pair sent by allowing an offset in the two-photon detuning for one probe-control pair (by adjusting the frequency sent to the control field AOM).

Fig. 11.1 has the timings for inputs/outputs and control fields for these protocols.

11.2 Experimental results

We performed a number of experiments on different occasions. In one experiment we focussed on maximising the efficiency of the memory and measuring the storage lifetime, while in another we investigated the sensitivity of the memory to various phase-matching effects.

11.2.1 Efficiency

Efficiency was measured with the assumption that a memory based on this technique would have the capability to send inputs with the correct relative phase for constructive interference into the spinwave. This made it necessary to estimate from the random phase data what the efficiency would be with the correct phase. The raw data from two runs, one with good phase for storage and one with bad, are shown in Fig. 11.2.

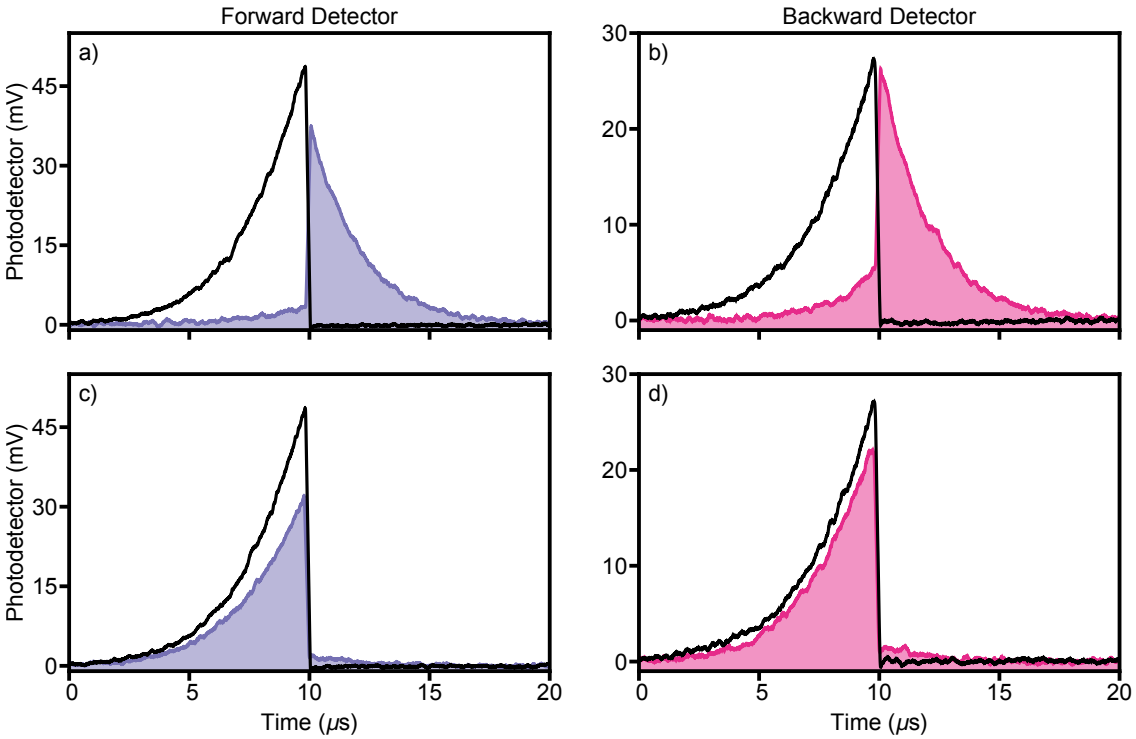

Fig. 11.2. Photodetector data for two different experimental runs. A reference pulse is sent with no MOT toward each photodetector (black). In one run (a,b), the input pulses (forward, blue and backward, purple) are efficiently absorbed. In the other run (c,d) the pulses are not absorbed due to destructive interference into the memory.

The efficiency was measured in two ways. First, the integrated output pulses were divided by their respective integrated reference pulse and then averaged. These are shown in Fig. 11.3. The distribution is consistent with a sinusoidal dependence on a random phase, where the maximal and minimal values are more likely than the intermediate values. The amount of noise present in the distribution was approximated by simulating a similar probability distribution in Mathematica. Based on this approximation, the 95th percentile of the data was taken to give a measure of the optimum efficiency.

The second method takes into account the relation between input leakage and output

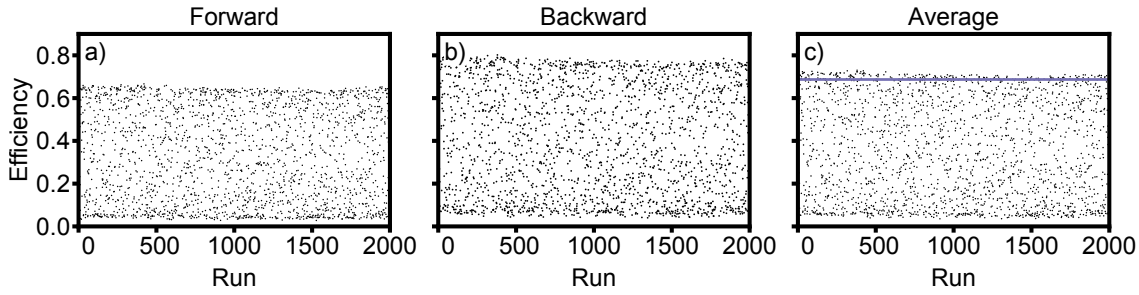

Fig. 11.3. Output pulse energy as a fraction of reference pulse received at detector. 10% of the averaged data are above the blue line.

recalled energy. For perfect phase-matching and coupling this would correspond to a line with the best efficiency at the lowest input intensity, or at the y-intercept. In practice, the phase-matching was not perfect, and the relationship produced an ellipse, as shown in Fig. 11.4. This ellipse of points was fit with least-squares according to the equation

$$a + bx + cx^2 + dy + ey^2 + fxy = 0. \tag{11.1}$$

The highest point of the fitted ellipse was used to estimate the efficiency that could be achieved if the relative phase of the couplings could be controlled exactly (but not the phase-matching).

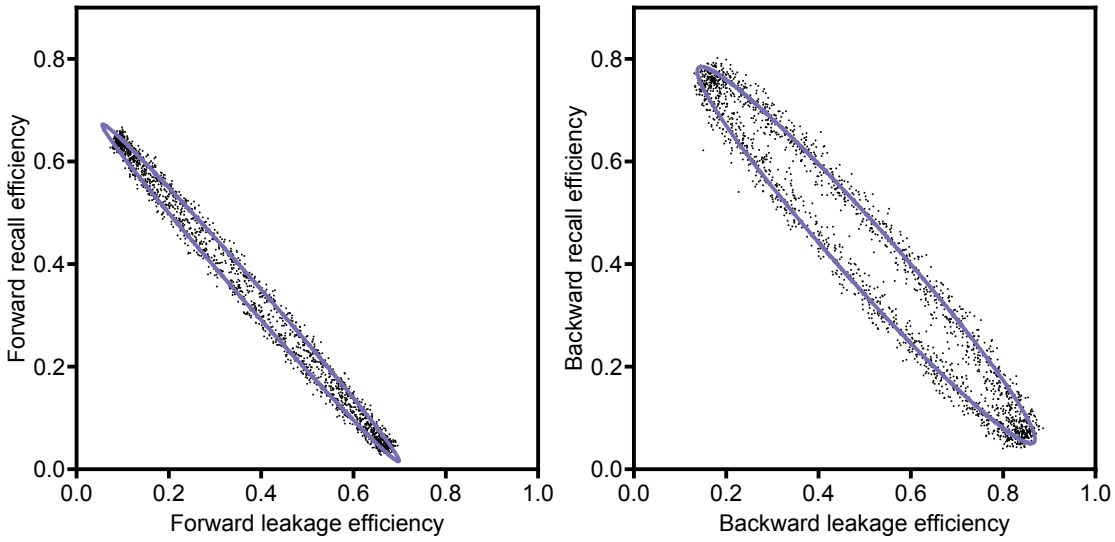

Fig. 11.4. Output pulse energy as a fraction of reference pulse received at detector as a function of input pulse energy leakage received. The fitted ellipses tend to overestimate the maximum recall, especially for lower eccentricity.

Efficiencies of $70 \pm 5\%$ were measured with the first dataset. The large uncertainty in the efficiency measurement is due partly to the uncertainty in the calibration of the detectors, and partly due to the imperfect methods of estimating efficiency from the data.

11.2.2 Phase-matching measurements

The relative phase between each input pulse in the first data set was essentially random. However, over a few hundred microseconds of a single experimental run, the relative phase is stable. The path lengths of the beams do not change significantly, and the experimental fields are all supplied from the same laser. This allowed us to send multiple pulses while changing the phase of the carrier wave of one of the probes. The measured data from a single run is shown in Fig. 11.5.

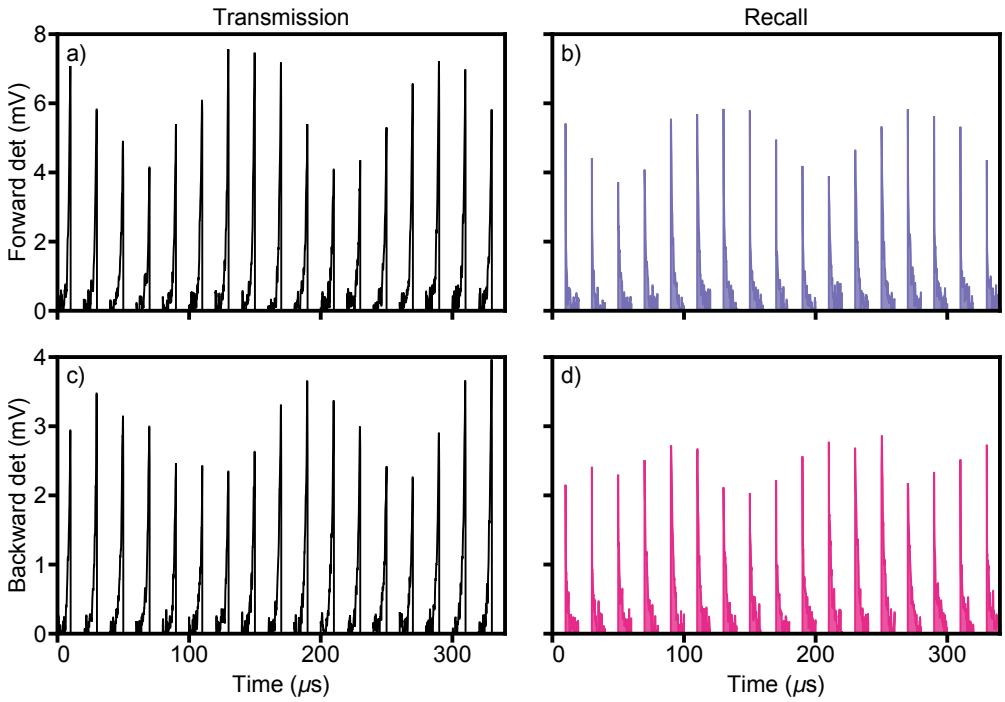

Fig. 11.5. Photodetector data from the phase increment measurements. Divided into input pulse transmission a) forward and c) backward, and recall b) forward and d) backward.

By changing the relative phase, we were able to measure more directly the phase dependence of the storage upon the two probes. The performance only depends on the relative phase of the pairs of fields, rather than any absolute phase, so from each experimental run a similar sinusoidal dependence could be extracted. This allowed a measure of the phase-matching, as well as an improved estimate of the maximum efficiency in a phase-controlled memory.

By plotting the various phase offsets against the input phase offset, the essentially constant relative phase between the various inputs and outputs could be determined. These are shown as ϕ_{OFF}, the intercepts of the linear fits in Fig. 11.7. This offset was constant within an experiment but was not reproducible due to its dependence on alignment. The top two rows of Fig. 11.7 illustrate how much the phase offset can change due to realigning the beams.

The optimal efficiency achieved for the phase-matching data was much lower than the dataset for efficiency measurement. These data were taken on a different day with a focus on measuring phase offsets for multiple datasets instead of optimising efficiency.

The effect on the phase matching of the additional excited state was observed by

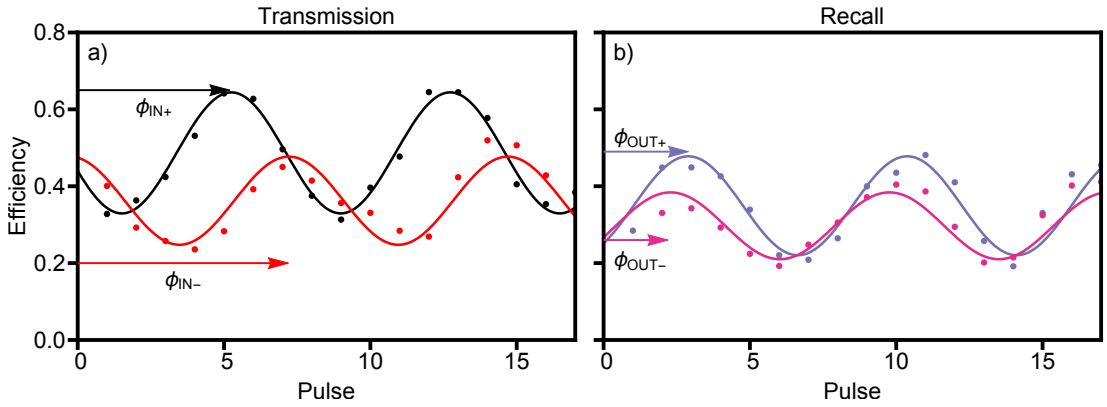

Fig. 11.6. Integrated pulse energies normalised to reference pulse at each detector. a) The transmission of the forward (black) and backward (red) input pulses. b) Recall of the spinwave at the forward (blue) and backward (purple) detector. The solid lines are sinusoidal fits to the data points. ϕ_{IN+} is the phase offset of the sinusoidal fit to the forward input data, and so on.

changing the one-photon detunings $\Delta_{\pm}$ by 20 MHz and measuring the corresponding shift in the phase offsets between inputs and outputs. This shift is shown as the differences in phase offset between the bottom two rows of Fig. 11.7. These measurements were consistent with the expected phase offset due to the dispersion. The phase offsets of the recalled outputs shift less than the phase offsets of the transmitted inputs.

This is due to the way in which the spinwave is driven. Upon storage, it is driven by a combination of the input fields and the field generated by the spinwave plus the control field. Upon retrieval, all the fields are generated within the memory. There is a phase offset at retrieval because a single directional input field will still store without a counterpropagating input. This storage is less efficient than the storage due to the interference of counterpropagating inputs and the phase offset is reduced at recall. This effect can also be seen in Fig. 11.8, where the recalled outputs have a smaller phase offset compared to the transmitted inputs.

11.2.3 Storage time

By switching off the control fields, the retrieval could be delayed. The decay time and any additional unexpected behaviour of the memory could be measured.

The decay data in Fig. 11.9 are fitted with an exponential function in time multiplied by a Gaussian, giving the function $\eta \propto \exp(-5.1 \times 10^3 t) \exp(-(4.9 \times 10^3 t)^2)$. The Gaussian decay envelope is associated with temperature, where the atomic motion due to the non-zero temperature causes the finer spatial frequencies of the spinwave to be washed out. Exponential decays are generally due to stray magnetic fields causing dephasing of the spinwave.

The decay profile of the memory was somewhere between Gaussian and exponential in time. The decay constant is consistent with a temperature of $120 \pm 50\mu$K. The large uncertainty in temperature is due to the fitting with exponential multiplied by Gaussian, which is highly dependent on the time window over which the data are fitted. The fitting could be expected to change if data were taken at longer storage times. Also, exponential decay might dominate over longer time periods where the transverse magnetic fields

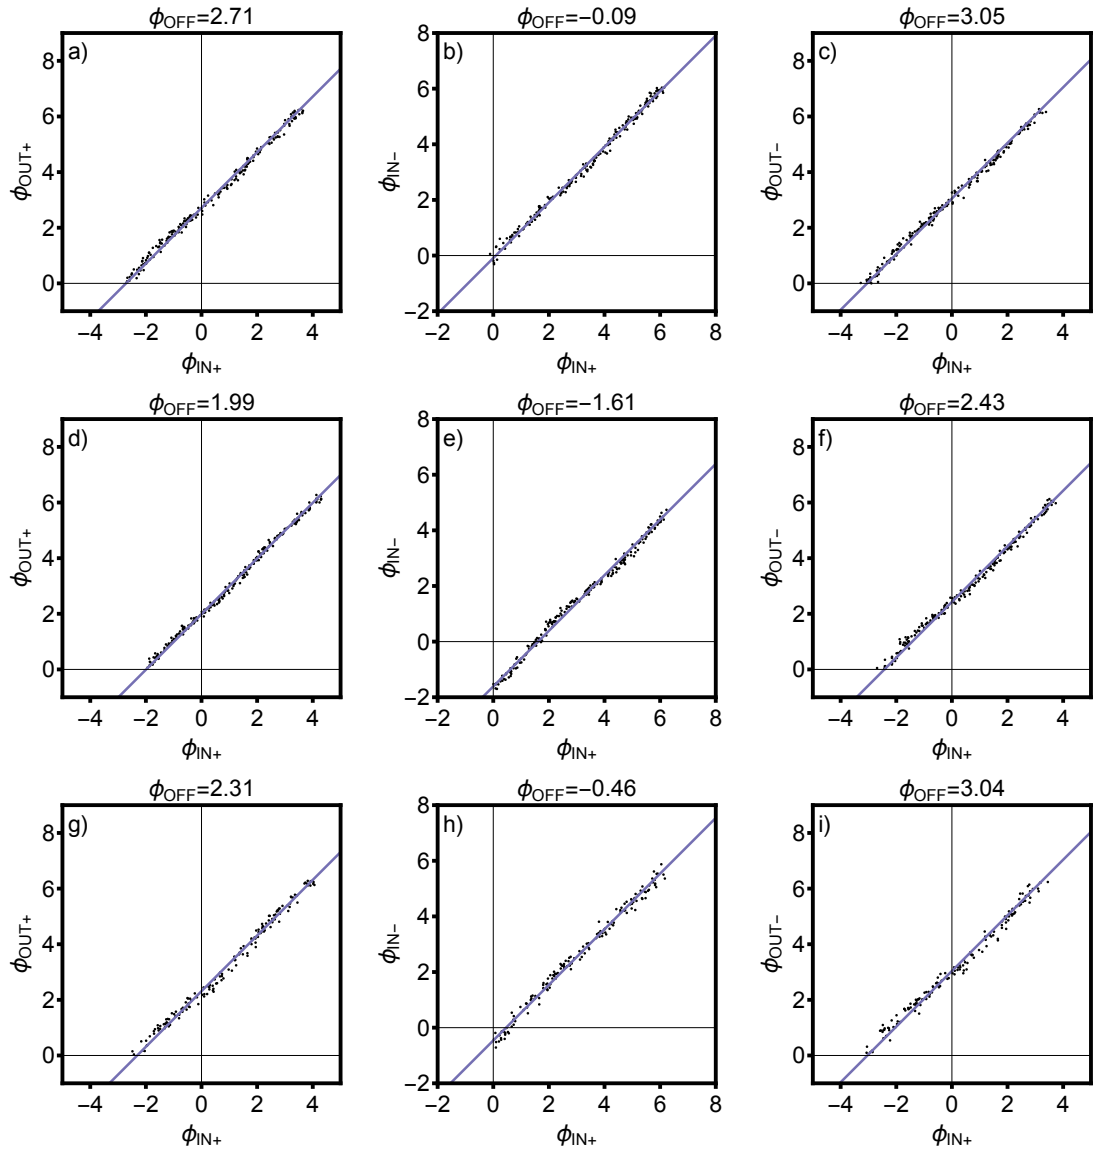

Fig. 11.7. Phases of the sinusoidal fits of all data as per Fig. 11.6 plotted against the forward input phase. The top row a-c) shows the phase offset between the forward input pulse and each of the other pulses for a well phase-matched experimental run. This can be seen as ϕ_{OFF} between the two inputs is only 0.09 rad and for the outputs is 0.34 rad (this second offset is not plotted directly but can be inferred from the relative phase of each to the forward input). The second row d-f) is a less well matched run, and g-i) shows the change in offsets when the alignment from the experiment for the second row data is preserved and the one-photon detuning is changed by 20 MHz.

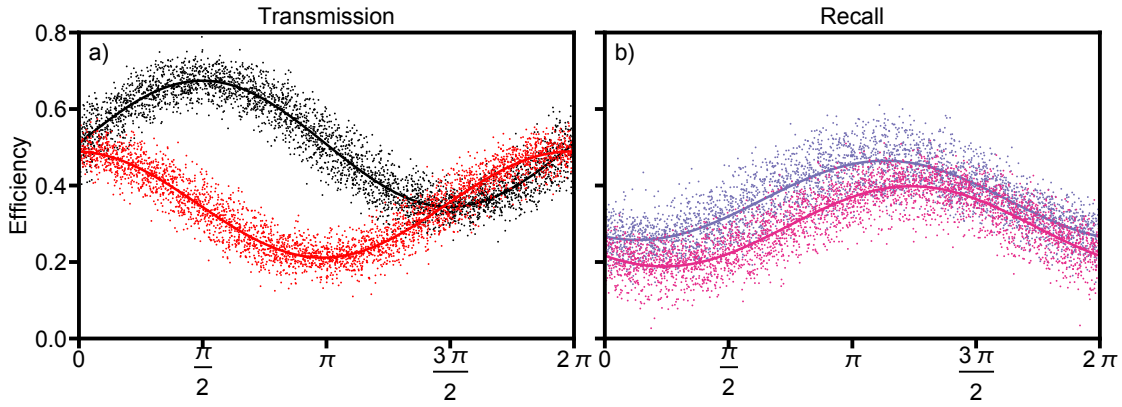

Fig. 11.8. Integrated pulse energies normalised to reference pulse at each detector, corresponding to the phase data in the second row of Fig. 11.7. The phase from each sinusoidal fit, relative to the forward transmission fit, is added to each data point to place it on the sinusoid. a) Forward (black) and backward (red) transmission. b) Forward (blue) and backward (purple) recall. Solid lines are sinusoids set to the average offset phase and then fit to the pulse amplitudes.

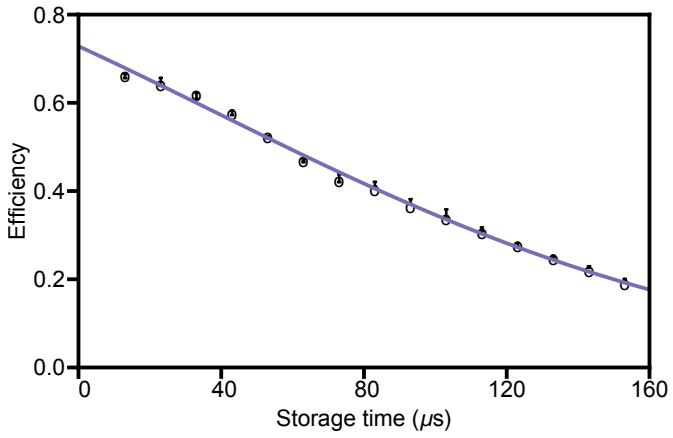

Fig. 11.9. Retrieval efficiency against storage time for exponential rising pulse input. Circles are maxima of fitted ellipses, error bars are the standard deviation of the largest 10% of outputs for each time. Line is a least-squares fit using the model of an exponential multiplied by Gaussian decay.

were not so effectively cancelled (cancellation is optimised at the start of the experimental run time window). This will tend to reduce the Gaussian component of the fitting, underestimating temperature-induced decay.

There are small regular oscillations in efficiency with time in the storage time data. The efficiency appears to oscillate above and below the fitted curve with a regular periodicity in storage time. There may be some minor inhomogeneity in the magnetic field across the memory, or in the behaviour of the second level on which the coherence exists. The TRACE memory is more sensitive to stray magnetic gradients than a gradient-echo memory. Small gradients in the longitudinal direction (which is the longest dimension and thus the most sensitive to stray fields) merely change the recall time and output bandwidth of GEM to a small degree. The gradients off the longitudinal axis produce a small change in the output direction, and introduce small pulse envelope distortions. For TRACE, these gradients reduce the recalled intensity, as there is no rephasing of the spinwave. The gradient may also change the effect of the phase mismatch, increasing or decreasing the recall efficiency.

11.3 Future experimental improvements

The setup used in this experimental demonstration of the free-space single-mode memory limited the characterisation of the memory, as well as the achievable efficiency. There are some improvements to the scheme that would make a more reliable experiment, and a more reliable quantum memory.

In a practical setting, the state desired to be stored will generally be in the form of a single optical pulse. The reliance on symmetric detuning for dispersion compensation makes it necessary to split the input pulse and shift the frequency of at least one part. This will inevitably add a significant amount of loss into the experiment. Instead, depending on the level structure of the ensemble used in the memory, it is possible to find optical frequencies at which there is no dispersion, as the dispersion is cancelled by the presence of transitions with both positive and negative detuning. The strengths of the Raman transitions do not generally cancel to zero as the transition strengths for each control field transition are generally different. A zero-dispersion transition then has the advantage of only requiring one probe and one control frequency, removing the need for frequency-shifting optics in the probe beam. This can possibly allow simpler geometries. A zero-dispersion transition will also not distort the spatial mode of the probe field and will make mode-matching more straightforward.

A single-frequency counter-propagating scheme requires consideration of the higher-order coherences in the spinwave generated by the direct interaction of probe fields with counter-propagating control fields. Where diffusion of the spinwave due to atomic motion is sufficiently fast, a single frequency can be used.

Passive phase-matching is possible in the Sagnac configuration of Fig. 11.10. The relative phase between each co-propagating pair of fields at the ensemble depends on the arm length, and will cycle through a full 2π radians for a path-length difference of the spinwave wavelength, $\Delta L = \lambda_{SW} = c/(\omega_p - \omega_c)$. The path-length differences within the pairs due to separating the control fields can be accounted for with either a fine path length adjustment of a separated control field, a coarse adjustment in the arm length where the fields are still combined, or a change in the ellipticity of the polarisation of a field before the input beamsplitter.

The output polarisation can be adjusted by changing the polarisation of the control field at recall. There will always be equal V and H components, but their relative phase can change the polarisation between two orthogonal polarisations in the basis at 45 degrees to V and H. This could be used to control the path of the recalled probe pulse through subsequent parts of the system without directly controlling the polarisation of the probe.

Generating the two inputs from a single-frequency pulse makes a Sagnac geometry more practical; Passive phase-matching of the control and probe fields requires each field in the pair to have similar paths through the interferometer, particularly bouncing off the same optics. Active optics such as acousto-optic modulators add complexity. The loss due to using symmetric detuning to reach a higher efficiency (the maximum efficiency is lower on the single-frequency scheme due to the reduced effective optical depth) is less than any expected loss due to frequency shifting. If an angle between control and probe fields is not required, a cavity or other frequency-discriminating optic can be used to combine or split the probe and control field.

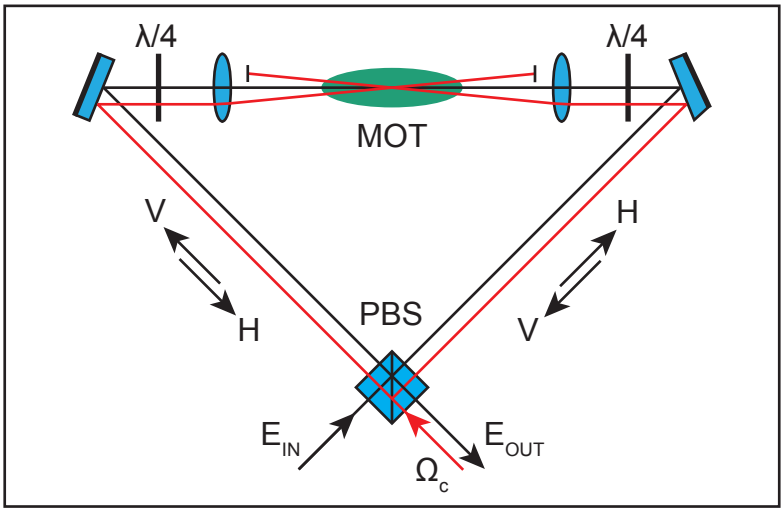

Fig. 11.10. Schematic for a single-mode memory with a Sagnac interferometer geometry. The PBS combines the probe and control fields with orthogonal polarisation between control-probe pairs. The quarter-wave plates set the circular polarisation at the MOT so both pairs of fields interact with the same transition. This changes the polarisation from H to V across the two quarter-wave plates, and the input and output ports are different. The scheme can be modified for same polarisation between control-probe pairs and/or no angular phase-matching by moving the input control field. The polarisation of the control field determines the relative phase between the two output fields and will determine the output polarisation.

11.4 Conclusion

In this chapter, we demonstrated an efficient single-mode memory in free space in a high optical depth ensemble of rubidium-87. We showed how the optimal input mode under constant control field is exponentially rising with an exponentially falling output resulting from time reversal. We found that phase-matching is a more important issue for this experiment than in our implementation of Raman-SL, possibly due to the higher optical

depth and the use of two counter-propagating inputs instead of one. We also showed how a more practical memory might be built using the Sagnac interferometer geometry.

Part III

Nonlinear atom-optics with Raman stationary light

Chapter 12

Theory and literature review of optical cross-phase modulation in ensembles

> Human beings and fundamental particles share one absolute commonality: they exist in their interactions. In between times, their positions and trajectories are indecipherable even to themselves.
>
> Nick Harkaway, *Gnomon*

We arrive at the most exciting application for quantum atom optics: The prospect of building a quantum computer that operates on photonic qubits. In this version, large entangled states of single photons are prepared via one and two qubit gates. This allows universal quantum computation, which includes all the advantages and possibilities associated with quantum computing. Apart from good photon sources, high efficiency optics, faithful storage, and linear processing, this also requires for simpler implementations large nonlinear interactions between photons capable of generating a π phase shift of the optical carrier wave. This phase-shifting interaction is called cross-phase modulation (XPM).

In this chapter, we will discuss the theory behind generating nonlinear interactions between photons. We will cover some of the available interactions and methods for enhancing these interactions. We will also cover the technicalities of generating large phase shifts while maintaining high fidelity processes.

12.1 Theory of cross-phase modulation

We focus our theoretical attention on Kerr nonlinearities in atoms. This describes the class of interactions where the local electric field generated by one optical state affects the atom, which in turn affects the evolution of another state. For the case of XPM, single photon level states must produce a sufficient interaction to impart a large phase shift.

We will also briefly mention using Rydberg atoms for nonlinearities. While the Rydberg blockade mechanism itself is not so relevant, the methods used to modify the propagation of light are similar to those used with Kerr nonlinearities. Therefore, some of the same limitations, and the same ways of dealing with these limitations apply. In particular, the optical band gap can be used in similar ways to enhance the interaction.

12.1.1 Control-phase gate with single photons

We will take initial quantum states stored in a memory, that will have some spatial profile. These can be described according to the distribution of the state within the memory, (see [157] for a more complete calculation).

$$|1\rangle_A |0\rangle_B = \int \alpha(k)\hat{a}_A^\dagger(k) |0\rangle_A |0\rangle_B \, \mathrm{d}k, \tag{12.1}$$

For input states of the form

$$|\Psi_{IN}\rangle = |1\rangle_A |1\rangle_B , \tag{12.2}$$

a conditional phase shift will then reliably perform the mapping

$$|1\rangle_A |1\rangle_B \rightarrow e^{i\phi} |1\rangle_A |1\rangle_B \tag{12.3}$$

with ϕ ideally on the order of π. Phase shifts a few orders of magnitude lower may be sufficient [158], but schemes based on larger phase shifts are simpler.

The phase shift is conditional on each mode containing a photon. All other states should not receive a phase shift, eg. $|1\rangle_A |0\rangle_B \rightarrow |1\rangle_A |0\rangle_B$. Eq. (12.3) also requires that the phase shift should be uniform in time and other coordinates of the states, and should not produce any distortion of the amplitude envelope of the state. This is necessary for the linear operations that are used in addition to XPM for building gates.

The extent to which this requirement is satisfied can be measured by

$$\sqrt{F}e^{i\phi} = \langle \Psi_{IN} | \Psi_{OUT} \rangle = \int \mathrm{d}z_1 \int \mathrm{d}z_2 \psi_0^*(z_1, z_2, t)\psi(z_1, z_2, t), \tag{12.4}$$

where $\psi(z_1, z_2, t)$ is the two-particle operator describing the evolved state at time t and $\psi_0(z_1, z_2, t)$ is an equivalent state that has not undergone the phase shifting interaction [157].

Retaining the pure state representation of Eq. (12.3) also implies that no entanglement should be generated between other coordinates (such as frequency) of the state. This entanglement is necessarily traced out due to projective phase measurements imposed by a combination of the linear operations (such as interferometry, see Sec. 12.3 and the projective measurements involved in the computation).

12.1.2 Cross phase modulation via AC-Stark shift

A strong interaction between photons can only be generated in the presence of matter. We have seen in chapter 2 that the oscillating electric field of light shifts the energy levels of an atom to produce the AC-Stark shift [159], according to the equation

$$\Delta_{AC} = -\frac{\Omega_s^2 \Delta\Gamma}{\Delta^2 + \Gamma^2}, \tag{12.5}$$

where the AC-Stark shift Δ_{AC} is the steady state energy shift of the transition for an atom being driven by a signal pulse with local Rabi frequency Ω_s. For $\int_{-\infty}^{\infty} \left|\hat{\mathcal{E}}(t)\right|^2 \mathrm{d}t = N_p$,

where N_p is the number of photons in the signal pulse, the Rabi frequency of the signal pulse is

$$\Omega_s^2 = \frac{\Gamma\sigma_{at}}{A}\left|\hat{\mathcal{E}}(t)\right|^2, \tag{12.6}$$

with σ_{at} the interaction cross-section and A the transverse area of the field about the point.

This phase shift is typically very small. However at the maximum possible interaction strength, where the light is tightly focussed so that $A \approx \sigma_{at}$, and with the detuning $\Delta = \Gamma$, an integrated shift of order $\int \Delta_{AC} \mathrm{d}t = \pi$ can be generated for an atom in free space.

A photon at this close detuning has a high chance of being incoherently scattered by the atom, reducing the fidelity of the process. The assumption of a steady-state shift may not be appropriate for schemes where one or both photons are propagating through an ensemble. The detuning can be increased in order to reduce the scattering probability. Increasing the integrated intensity of the light by placing the interaction inside a cavity or enhancing the interaction with an additional atom-light interaction can make up for the increased detuning. This can then increase the maximum fidelity and success probability of the XPM process.

Much of the work we review is based on this principle; both as demonstrations of the enhancement, and further refinements and limitations on the effectiveness of the various schemes.

12.1.3 Dispersion of optical fields by atoms in free space

The AC-Stark effect is necessarily associated with polarisation of the atoms by the modulating field. The coherent re-emission by the atoms produces a new optical field that is shifted in phase by $\approx \pi$. The overall phase shift will depend on how much of the field is coherently absorbed and re-emitted. This depends in a complex way on the parameters of the field [162]. That work includes calculation of the field emitted by an atom with an incident weak coherent optical field, with a well-defined electric field amplitude at the atom. By calculating the dipole radiation of the atom, the phase shift of the outgoing light is then calculated.

The dispersion of the optical field can be related to the polarisability. We have from (12.5) that the AC-Stark shift is proportional to the susceptibility multiplied by intensity.

$$\Delta_{AC} \propto \left|\hat{\mathcal{E}}\right|^2 \mathrm{Re}[\chi] \tag{12.7}$$

In the steady state picture, as long as the atom is far below saturation, the dispersion should not depend on the strength of the field. Neither should the integrated AC-Stark shift of the atom depend on the time profile of the photon, as long as the bandwidth does not affect the detuning from the transition Δ. This allows for the intuitive linking of the two phenomena - dispersion and AC-Stark shift. Inducing a phase shift on the atom is equivalent to a measurement of the existence of the photon, and must therefore produce some effect on the photon. Provided the interaction is carefully designed, this generates entanglement rather than a decoherence-inducing measurement.

The relation allows for some useful conclusions. Provided the atom is not saturated, the phase shift induced on the atom and the dispersion of the photon only depend on the focussing of the photon in space, and not the compression of the photon in time. Therefore, as evident from the AC-Stark shift calculation, any enhancement of the interaction

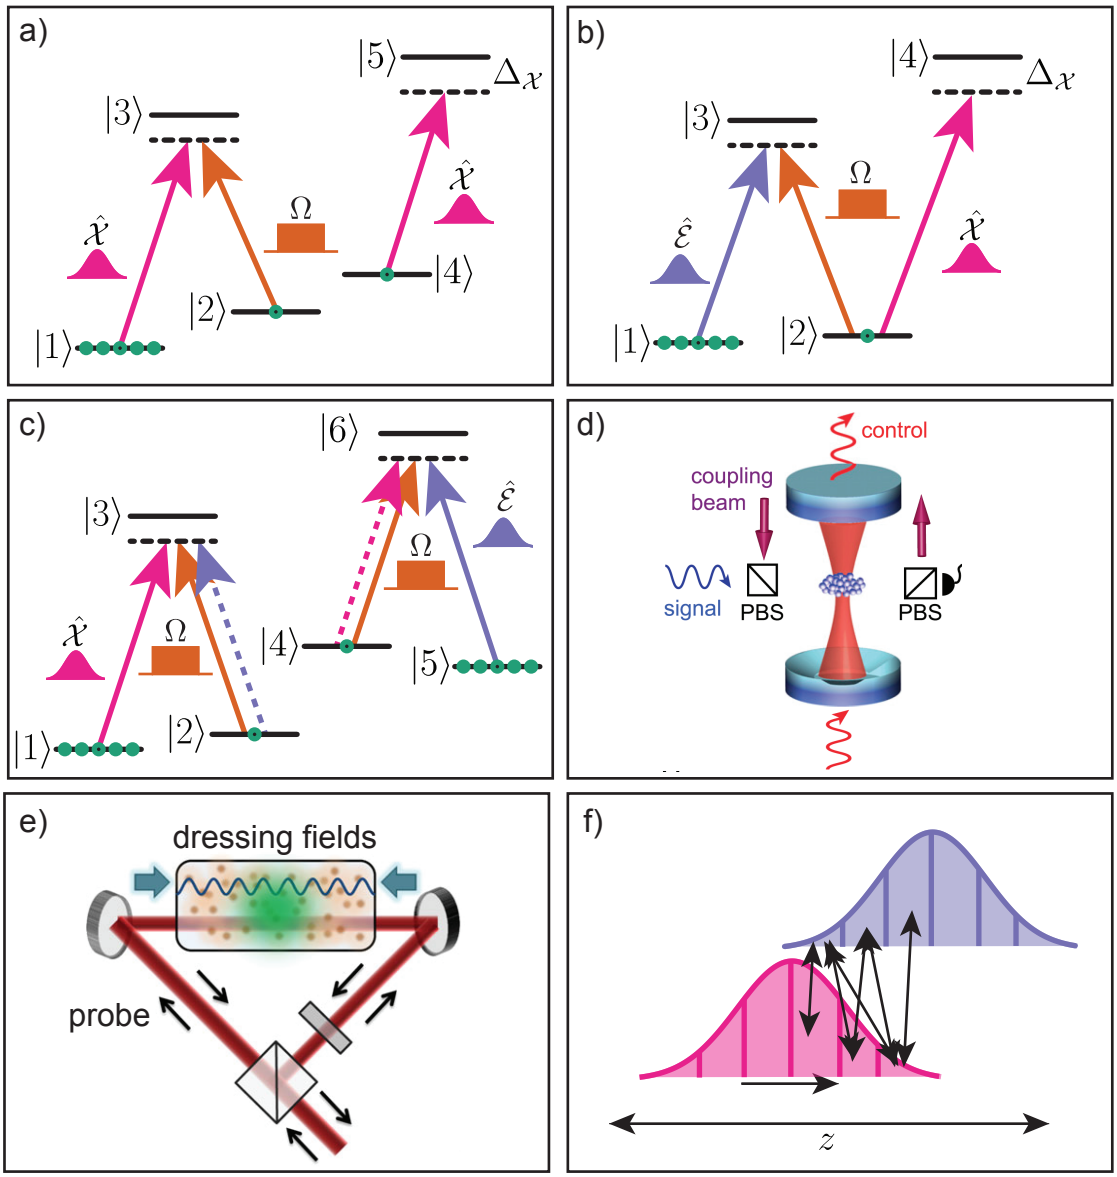

Fig. 12.1. Schemes for ensemble enhanced XPM. a) the ensemble modifies or enhances the modulating field $\hat{\mathcal{X}}$ via the interaction with levels $|1-3\rangle$ and a control field Ω, while the field also interacts with a second coherence. The cross-phase shift is generated by the AC-Stark shift of $|4\rangle$. b) the field $\hat{\mathcal{X}}$ modifies the interaction of the probe $\hat{\mathcal{E}}$ with the ensemble by changing the detuning of the control field by AC-Stark shifting level $|4\rangle$. In the case of EIT, changing the two-photon detuning changes the dispersion experienced by the probe, causing a phase shift of the transmitted light. c) two EIT pulses each modulate the other . d) From [160] the 'control' affects the phase of the signal photon stored with EIT in the ensemble. The cavity increases the intensity of the control. e) from [161] the coupling of the probe is spatially modulated by a standing wave in the dressing fields, resulting in reflection throughout the ensemble. This increases the intensity of the probe. f) Entanglement forms between different components of the pulse as they cross. If the pulses do not cross completely or if the nonlinear interaction modifies the propagation of the pulses, these entanglements may not 'average out'. The entanglement then adds phase noise.

must involve an increase in the integrated intensity that the atom experiences.

Considering the phase shift of the optical state in terms of the emission from the modulated atom also helps explain the situation of an ensemble with a spinwave composed of a single spin (a single absorbed photon). All the emitters are in a superposition of being in the interacting level, and so emit when the modulating field is applied. The emitted mode is then the coherent superposition of that of the individual emitters. The overall dispersion is still roughly equivalent to that of a single atom, although the mode may be affected by the spinwave distribution. The ensemble also evolves according to the AC-Stark shift at each atomic location. There is therefore a straightforward extension of the above theory to ensembles, with the caveat that ensemble storage may require larger volumes than single atom storage, limiting the focussing and thus the available cross-phase shift.

12.2 Literature review on cross-phase modulation with atomic ensembles

There are many platforms for generating large cross-phase shifts between quantum states. We will focus on the use of atomic ensembles in combination with optical quantum states. We will address many of the proposed schemes and experimental demonstrations, as well as several no-go theorems that point out the limitations of many interactions for producing high fidelity cross-phase shifts. This will allow a thorough treatment of all the literature relevant to our exploration of using Raman stationary light to enhance XPM in Chapter 13.

12.2.1 AC-Stark interaction

In these schemes, no EIT is necessary. Instead, the modulating field simply shifts the frequency of a transition that the probe is near resonant with, or stored upon. The change in dispersion or output phase conditional on the presence of the modulating field produces the conditional phase shift. This is illustrated in the interaction between the modulating field $\hat{\mathcal{X}}$ and levels $|4\rangle \rightarrow |5\rangle$ in Fig. 12.1 a), ignoring the EIT scheme on the other levels.

For example, by confining the light in a hollow core photonic fibre, [163] was able to generate 0.3 mrad per photon. Modulation of a state stored by AFC in a waveguide [164] produced 10^{-9} rad per photon.

12.2.2 Cavity enhancement of cross-phase modulation

The intensity is increased by use of a cavity. The cavity multiplies the number of traversals and so the total phase shift. Generally, one state is stored so there are no additional complications due to co-propagation (this was a limiting factor in an experiment by Turchette et al. in 1995 [165], where the photons required bandwidths less than the cavity bandwidth to be spectrally separate, limiting the interaction time below the photon duration). The challenge for cavity enhancement is to place a high OD ensemble of atoms within the cavity mode to allow for efficient storage while still being able to use a sufficiently high finesse cavity.

In Beck et al. in 2016 [160], a weak control beam resonant with a cavity of cooperativity 3.8 was used to impart a phase shift on a probe photon stored via EIT in the ensemble. This was purely an AC-Stark scheme - the EIT coupling was used to reliably store the

probe and switched off during the XPM stage. A phase shift conditioned on the detection of a control photon of $\pi/6$ was measured. The experimental layout is shown in Fig. 12.1 d).

The use of high finesse cavities with single atoms has also been widely studied. For example EIT was explored with one or two atoms in a cavity [166]. A whispering gallery resonator has been used to both strongly localise the probe field and increase the number of traversals, achieving a phase shift per photon close to π [167]. These interactions are qualitatively different to ensemble effects due to the small number of emitters.

12.2.3 Enhancement of cross-phase modulation by EIT

A scheme equivalent to Fig. 12.1 b) was proposed, predicting orders of magnitude enhancement in XPM over a similar scheme absent EIT [168]. The AC-Stark shift of the EIT bandwidth changes the dispersion of the traversing photon, resulting in a phase shift. This shift is dependent on the optical depth of the ensemble, as this determines the relation between the width of the EIT window and the subsequent phase shift, and the group velocity of the probe, and thus the total interaction time. This interaction was in fact indirectly measured in the landmark EIT paper [20] due to interactions between the control field and another atomic level. This scheme was analysed further to find the exact phase shift expected for the ensemble parameters [169].

The scheme can also be used for single photon switching, where one photon is transmitted but two photons are not [170]. This was demonstrated in a rubidium MOT [171], where a laser pulse containing 23 photons per atomic cross-section caused a $1/e$ absorption of the second pulse. In another demonstration a hollow core photonic fibre was used to produce an attenuation of 3 dB for a single photon per atomic cross-section [172].

There have been several demonstrations of XPM based on this scheme [173, 174]. [175] uses a double EIT scheme to send two co-propagating pulses, measuring a shift of 1 μrad per photon. Phase shifts of order π but with high loss are shown with small coherent pulses in a similar scheme in [176].

The most impressive demonstration for single photons to date is from 2015 [177], where a phase shift of 18 μrad for post-selected single photons (sent as a weak coherent pulse) was observed.

The analysis of EIT outside the steady-state assumption also has implications for XPM [178, 179]. The finite rise time of the EIT transparency may improve the possible cross-phase shift [180, 181, 182]. An experimental demonstration of EIT transients was discussed [183]. The finite response time of the EIT modulation is shown to enhance the cross-phase shift even where the EIT window has narrower bandwidth than the pulse bandwidth [184]. The finite response time of the medium due to the narrow bandwidth extends the time over which the nonlinear interaction occurs. This does not take into account the loss of the photons due to the bandwidth mismatch. Again, this finds an enhancement of the integrated XPS proportional to optical depth, and the scheme requires the states to be co-propagating optical states. This effect has also been demonstrated experimentally [185].

12.2.4 Enhancement of cross-phase modulation by gradient echo memory

The propagation of light stored in the gradient echo memory and again placed under a control field is capable of producing enhancement of XPM [186]. A shift of order 10^{-12}

rad per photon was demonstrated, with higher phase shifts possible in principle. The proposal used a scheme similar to that of Fig. 12.1 a), with the control field reproducing the modulating field $\hat{\mathcal{X}}$ which then AC-Stark shifted a second coherence.

12.2.5 Enhancement of cross-phase modulation by stationary light

Stationary light was originally proposed with enhancement of nonlinearities in mind [116], and a method for doing so can also be found in Andre et al. (2005) [119]. There are yet to be any experimental demonstrations of XPM in stationary light, and there has not been much theoretical attention to the subject either. This sparsity is a little unsatisfactory, though much of the theory for EIT enhancement applies to EIT-SL. In particular some of the concerns raised in Sec. 12.2.8 regarding phase noise due to unwanted entanglement apply to EIT-SL. In Chapter 13, although we focus on Raman-SL, we do compare the expected enhancement to EIT and EIT-SL.

In terms of demonstrations, a scheme similar to the co-propagating scheme of [168] was proposed [187]. Here, the modulating field was the optical field of a state trapped as stationary light.

An analogue of the photonic switch was demonstrated [188]. The erasure of one spin-wave by a stationary light pulse was used to indicate the presence and intensity of the stationary light.

12.2.6 Rydberg blockade enhancement of cross-phase modulation

The direct state-dependent interaction between atoms is generally negligible in cold vapour atomic ensembles. However, atoms with an electron in an orbital state with high principal quantum number n (Rydberg atoms) have a strong dipole-dipole interaction. The presence of a Rydberg atom significantly shifts the energy levels of the Rydberg states for nearby atoms. That shift can prevent the creation of another Rydberg interaction within a small distance, in an effect known as Rydberg blockade. The interaction can be used to generate a conditional phase shift or a number of other nonlinear optical interactions [189]

A Rydberg blockade in combination with EIT has been proposed for generating conditional cross-phase shifts [190, 191]. The EIT is generated by a control field coupling the excited state to a Rydberg state. The ensemble nature of the EIT interaction enhances the Rydberg blockade effect: The blockade modifies the transparency effect, affecting the propagation of polaritons passing through the blockade and imprinting a phase shift. This is equivalent to the way in which the AC-Stark shift interacts with EIT for enhanced XPM, but the interaction per photon via Rydberg blockade can be significantly stronger, while nonlocality of the interaction can allow more sophisticated schemes.

The effective optical depth for enhancement of Rydberg interactions is the optical depth in the Rydberg blockade radius. The optical depth enhances the interaction by allowing the superposition of many contributors to the Rydberg interaction (just as many emitters increase the optical field for the AC-Stark shift, or decrease the group velocity of the probe under EIT). This can be overcome by placing the ensemble in a cavity, where the multiple passes of the probe through the ensemble increase the interaction without introducing multi-mode entanglement [192, 193, 194].

12.2.7 Distributed reflection or induced cavity

The investigation of stationary light is also motivated by other modifications of light propagation, not just the increased intensity time-integral. The generation of a distributed reflection can be used to enhance nonlinear interactions, while avoiding the problems associated with such enhancements in a solely co-propagating case. [195].

Working in the context of Rydberg atoms, Lahad et al. proposed an 'induced cavity' where the spatial grating is induced in the nonlinear interaction itself [161]. The system is then operated outside the bandgap in a transmission maximum as in [150]. This method produces intensity enhancement due to the reflection throughout the medium, without the dynamic control necessary to stop a pulse and convert it to stationary light. The output characteristics of the light are then similar to those for a cavity. The geometry for the proposed scheme is shown in Fig. 12.1 e).

12.2.8 Limitations of ensemble cross-phase modulation and no-go theorems

Study of the multi-mode nature of cross-phase modulation interactions has resulted in several 'no-go' theorems. That is, it was shown that the enhancements originally proposed were insufficient to produce high fidelity large conditional phase shifts. These do not rule out the possibility of XPM. Rather, they impose additional requirements that increase the difficulty of doing so. They also imply that any proposed schemes should be examined thoroughly for the sorts of problems found with previous schemes, and place more onus on experimental demonstration to show that the claims are warranted.

Nonlinear interactions must be designed to avoid entanglement between properties other than the intended degree of freedom. Any unwanted entanglement will be 'traced out' upon measurement and produce a mixed state contribution to the final state. This lowers the fidelity of the operation.

The impossibility of improving Kerr-type nonlinearities between co-propagating modes by only increasing propagation distance is a well-known example [196]. The local causal nature of the interaction causes problems including re-emission of the target mode into empty modes. There is a necessary compromise between fidelity and phase shift.

The argument is applied by Shapiro and Razavi [197] to show the limitations on cross-phase shift, by examining the phase noise due to spectral entanglement. This is interesting in that it takes the rise time of the XPM into account to place the limitation on both fast and slow interactions. The authors also suggests a full field-operator treatment of the process is necessary to discover the exact relation between phase shift. An example of spectral entanglement arising is shown in Fig. 12.1 f).

EIT enhancement is examined specifically by Gea-Banacloche [198]. The relationship between dispersion and EIT transparency window means that any large phase shift is associated with a large absorption of the probe; where the window is wider than the pulse bandwidth, there is little dispersion and no absorption. Where the window is narrower, there is large dispersion and absorption.

A continuous-mode picture of EIT enhanced XPM is studied by He et al. [157], detailing the dependence of the fidelity of the interaction on the optical depth and the phase shift. A uniform phase shift is not guaranteed even for counter-propagating pulses, and again the finite bandwidth of the interaction is a limiting factor for the fidelity. He et al. explain that a coordinate error in some previous works was responsible for overestimating the possible fidelity for large phase shift.

[199] demonstrates numerically the effect of co- versus counter-propagating pulses by discretising the interaction into a number of sites, while applying a multi-mode treatment of the photons. It is shown in [200] that only 12 sites are required (for a one-dimensional waveguide) for a fidelity $\mathcal{F} > 0.99$.

The transverse distribution of the optical fields is shown to have an adverse effect on the fidelity [201]. The interaction strength, and the resulting phase shift varies according to the local amplitude. For nonlocal interactions such as Rydberg blockade, this can be avoided by a sufficient transverse separation of the pulses. For local interactions, strong transverse confinement such as by a waveguide may be a solution, provided that the waveguide mode can coexist with the nonlinear interaction.

So far, there has been no direct experimental comparison of co- and counter-propagating interacting pulses. Instead, it has been schemes that use photons circulating in cavities that have shown the larger phase shifts. Confining the interacting photons to single modes by placing them in a cavity is expected to resolve many of the above issues with multi-mode entanglement.

For schemes where multi-mode entanglement may be a problem, demonstrations have started to address those concerns. For example, in [164] the independence of the phase shift on the exact time-bin mode of the modulating field is shown. The possibility of large high fidelity phase shifts might only be resolved by experimental demonstrations of conditional XPM creating entangled qubits, with full tomography to reveal any unexpected sources of noise.

12.3 Optical quantum computing

Quantum computing (QC) has emerged in the last few decades as a potential application of quantum information science. The roots of QC can be traced to Feynman's 1982 conjecture [202] that a quantum computer could more efficiently simulate quantum systems than a classical computer. The benefits of such simulations are anticipated today, with potential advances in material and biological sciences, as well as fundamental physics.

Understanding of the required algorithms to perform universal QC also relied on work in reversible computing [203, 204]. Quantum mechanics is inherently unitary and thus reversible. In particular, the interactions we are examining must be unitary to achieve high fidelity operations. The gate proposed in [204] allows universal computation, and its quantum analogue, the control-swap gate allows for universal QC [205].

Apart from quantum simulation, QC was shown to have advantages for more conventional applications. For example, Shor's algorithm [206] is famous for promising efficient factorisation of large numbers, threatening the security of modern encryption and the economic activity and global communication it supports. For a more benign application, Grover's algorithm [207] would allow for efficient database searches.

A combination of quantum communication with quantum computing to build a 'quantum internet' [208], would allow even more novel applications in terms of communication and computing. This is also a useful review of the diverse range of systems being investigated for the necessary interactions.

12.3.1 Optical quantum computing

We have mentioned optical quantum computing as the motivation for studying nonlinear atom-light interactions. We will briefly explore what these interactions would look like

in the context of optical QC in order to make sense of how the XPM interactions fit into the larger scheme. This will also point to how the interactions should be implemented.

As pointed out in [209], for weakly interacting qubits such as photonic qubits, generating the required interference effect is not difficult. For strongly interacting qubits, generating the required nonlinear interaction is not difficult. Generating both simultaneously is a challenge. As in Fig. 12.2, implementing the control-phase gate part would require going beyond even the impressive work reviewed in this chapter, while the interferometer part is standard equipment.

It is possible to build a universal quantum computer out of control-phase gates and linear operations. The implementation for a Fredkin or control-swap gate illustrated in 12.2 is an example of a universal gate that could operate on photonic qubits.

The control-phase gate can be implemented using, for example, polarisation qubits. A polarisation qubit might consist of a single photon in a well defined time and spatial mode, existing in a superposition state of two orthogonal polarisations. For example, a qubit in state $|0\rangle$ would correspond to a left-circularly polarised photon $|0\rangle \equiv |L\rangle$, and state $|1\rangle$ to a right-circularly polarised photon $|1\rangle \equiv |R\rangle$). These polarisation states are easy to manipulate experimentally, allowing implementation of single qubit gates. They also have the advantage that atomic systems can easily be used to generate polarisation dependent interactions.

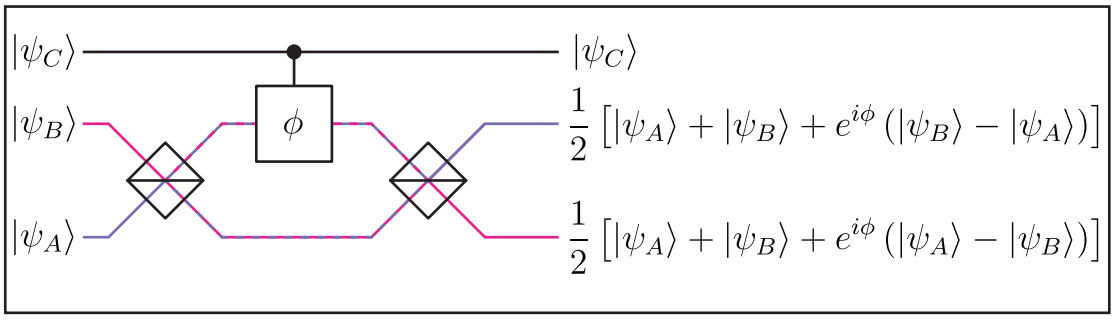

Fig. 12.2. The addition of the control-phase gate to an interferometer geometry allows for a control-swap or Fredkin gate [209]. For the phase shift $\phi = \pi$, the output maps the state in mode B to mode A, $|\psi_B\rangle_B \rightarrow |\psi_B\rangle_A$ and vice versa, $|\psi_A\rangle_A \rightarrow |\psi_A\rangle_B$. The swap is conditional on the presence of a photon in mode $|\psi_C\rangle_C$, thus generating entanglement between the states in modes A, B and C.

If the control qubit in Fig. 12.2, $|\psi_C\rangle_C$ is in the superposition $1/\sqrt{2}(|0\rangle_C + |1\rangle_C)$, the whole system will be mapped to the superposition (with entanglement between the control qubit and the other two qubits individually),

$$|\psi_A\rangle_A |\psi_C\rangle_C \rightarrow \frac{1}{\sqrt{2}} (|0\rangle_C \otimes |\psi_A\rangle_A + |1\rangle_C \otimes |\psi_B\rangle_B) \tag{12.8}$$

$$|\psi_B\rangle_B |\psi_C\rangle_C \rightarrow \frac{1}{\sqrt{2}} (|0\rangle_C \otimes |\psi_B\rangle_B + |1\rangle_C \otimes |\psi_A\rangle_B) \tag{12.9}$$

If the interaction is sensitive to the circular polarisation state, the qubits can be separated into their $|L\rangle$ and $|R\rangle$ components using polarising optics, sent along separate paths to allow uniform interactions, and recombined afterwards. For a control qubit, the sensitivity to polarisation may actually simplify the situation, as the whole qubit can be sent into the gate knowing that the interaction will still be conditional on the presence of the

$|R\rangle$ component.

Since the complete optical Fredkin gate requires fairly trivial additions to the control-phase gate (except using single photon qubit states instead of weak coherent states) there have been no demonstrations of a full gate operation.

However, Hacker et al. have demonstrated the generation of entanglement between the input photons that is possible with a control-phase gate [210]. This made use of a single atom in a high finesse cavity, where the phase of two subsequent reflected photons was dependent on the state of the atom. This was based on a protocol [211] where the atom in the cavity acts as an ancilla, allowing indirect interaction between the two photons.

12.4 Summary

Research in ensemble enhanced cross-phase modulation has made a great deal of progress recently, with phase shifts on the order of π demonstrated in several different systems. Enhancements with stationary light have not yet been performed, and the following chapter will explore the requirements for such a demonstration.

Chapter 13

Cross-phase modulation with Raman stationary light

> Consider what a great forest is set on fire by a small spark.
>
> James 3:5, New International Version

In this chapter, a cross-phase modulation (XPM) scheme based on Raman stationary light is explored. We are interested in what amount of cross-phase shift per photon could be expected. We also investigate how issues such as phase noise by unwanted entanglement might affect the scheme, and how to deal with this problem. A success rate is found for a simplified model of the XPM scheme. The experimental parameters such as optical depth and detuning necessary to build a reliable cross-phase gate based on the scheme are also found, and a scheme for implementing a controlled swap gate entirely within a memory is proposed.

The relevant published paper to this chapter is

"**Dynamical observations of self-stabilizing stationary light** J. L. Everett, G. T. Campbell, Y.-W. Cho, P. Vernaz-Gris, D. B. Higginbottom, O. Pinel, N. P. Robins, P. K. Lam, B. C. Buchler, Nature Physics 13(1), 68-73 (2016)"

13.1 Maximum cross-phase shift generated by a stationary light field

It was shown that a controlled phase gate in addition to linear operations is a sufficient set of operations for building a universal quantum computer [205]. We investigate using the AC-Stark shift to generate an interaction between an optical field and an ensemble state, and find the amount of phase shift that can be produced dependent on experimental parameters.

We make two assumptions to reduce this to a simpler problem. The first is to consider only the one-dimensional case. We assume that the interaction is not dependent on any transverse distribution of the states. This could be achieved either by designing the states so they are uniform in the transverse dimensions of the interaction region, or using a one-dimensional platform such as a waveguide that constrains the fields to a single transverse mode.

The second assumption is that the only source of loss is the scattering of the control field from the signal spinwave. The conditions for neglecting other sources of loss are explored in Sec. 8.4.

These assumptions allow the calculation of the XPM by the AC-Stark shift due to a stationary light field of a target state stored in the ensemble. The total cross-phase shift achievable in the scheme can be calculated by integrating the AC-Stark shift over time as the light decays. We assume that the stationary signal spinwave and the target state are each generated by the efficient storage of a single photon.

First, by calculating the amount of AC-Stark shift that is generated while the signal field decays, it is possible to show the dependence on optical depth.

The AC-Stark shift depends on the Rabi frequency and the detuning of the signal-field from the interacting atomic level. The level scheme is illustrated in Fig. 13.1

$$\Delta_{AC} = -\frac{|\Omega_s|^2}{4\delta_s} \tag{13.1}$$

We obtain a relation between the signal intensity and the Rabi frequency from the cross-phase shift produced by a free moving photon,

$$\theta_{XPS} = \int_{-\infty}^{\infty} \Delta_{AC} \mathrm{d}t = \frac{\Gamma \sigma_{at}}{4A\delta_s} \tag{13.2}$$

where σ_{at} is the effective interaction cross-section for the Stark-shifting transition, and A is the cross-sectional area of the signal field. For $\int_{-\infty}^{\infty} \left|\hat{\mathcal{E}}(t)\right|^2 \mathrm{d}t = N_s$, where N_s is the number of photons in the signal pulse,

$$\Omega_s^2 = \frac{\Gamma \sigma_{at}}{A} \left|\hat{\mathcal{E}}(t)\right|^2 \tag{13.3}$$

A signal pulse of the form $\hat{\mathcal{E}}_s(t) = \frac{1}{\sqrt{2\tau}} e^{-\frac{t^2}{4\tau^2}} (e^{i\omega_+ t} - e^{i\omega_- t})$ is stored in the memory over the time $-t_0$ to t_0. The sidebands $\omega_\pm$ are chosen such that the two frequencies are stored in separate halves of the memory and will generate stationary light. Assuming ideal GEM storage [73], we can approximate the spinwave as the Fourier transform

$$\hat{S}_s(z, t_0) = \frac{\eta\sqrt{\Gamma}}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp(iz\eta t) \hat{\mathcal{E}}_s(0, t) \mathrm{d}t \tag{13.4}$$

The transform is area preserving apart from a factor of $\sqrt{\Gamma}$:

$$\int_0^1 |\hat{S}_s(z, t_0)|^2 \mathrm{d}z = \Gamma \int |\hat{\mathcal{E}}_s(0, t)|^2 \mathrm{d}t \tag{13.5}$$

The two signal frequencies are absorbed into the separate halves of the memory, and integrating over half the memory gives

$$\int_0^{1/2} |\hat{S}_s(z, t_0)|^2 \mathrm{d}z = \frac{\Gamma}{2} \int |\hat{\mathcal{E}}_s(0, t)|^2 \mathrm{d}t \tag{13.6}$$

Using Eqs. (8.9), (8.10), (13.6) and integrating over half the memory to solve for Ω_s^2 at the

centre we find

$$\Omega_s^2 = \frac{\Gamma\sigma_{at}}{A}\int_0^{1/2}\left|\sqrt{d}\frac{\Omega_c}{\Delta_c}\hat{S}_s(z,t_0)\right|^2 \mathrm{d}z \times \exp(2\gamma(t_0-t)) \tag{13.7}$$

$$= N_s\frac{d\sigma_{at}\Gamma^2}{2A}\frac{\Omega_c^2}{\Delta_c^2}\exp\left(4\Gamma\frac{\Omega_c^2}{\Delta_c^2}(t_0-t)\right), \tag{13.8}$$

where Ω_c and Δ_c are the control field Rabi frequency and detuning respectively. The stationary light intensity and decay rates are both proportional to control-field intensity, so this factor cancels in the time integral and the total cross-phase shift generated in the scheme is

$$\phi_s = \int_{t_0}^{\infty}\Delta_{AC}\mathrm{d}t = -N_s\frac{\Gamma}{\delta_s}\frac{\sigma_{at}}{A}\frac{d}{32} \tag{13.9}$$

The amount of cross-phase shift is proportional to the optical depth, with the ratio of maximum cross-phase shift compared to a freely propagating photon equal to $d/8$. This scaling is common to other XPM schemes utilising the AC-Stark shift. However, the requirements for generating a high-fidelity gate can be different. This is because the optical field in the Raman stationary light scheme can exist outside the spinwave. In the EIT stationary light scheme, the optical field only exists in the same region as the spinwave.

Throwing away the one dimension assumption, to obtain a uniform phase shift across the target state, the intensity of the signal field must be uniform. Where the optical modes are Gaussian beams, uniformity requires that the target state must be stored along a region shorter than the Rayleigh length of the signal field. For EIT schemes, the relevant optical depth for enhancement is that of the interaction region. For the Raman stationary light scheme, the optical depth across the entire coherence determines the strength of the optical field in the interaction region.

The intensity of an optical field at the waist is inversely proportional to the area of the waist. The area is proportional to the Rayleigh length. If the beam can be focussed so the Rayleigh length is less than the length of the medium, tighter focussing decreases the Rayleigh length further, decreasing the available optical depth. Therefore it is the density of the medium, rather than the total optical depth, which determines the amount of phase shift possible with an EIT scheme. A longer ensemble at the same density may be beneficial for the Raman stationary light scheme, assuming the issue of defocussing (see Sec. 8.5.1) can be dealt with.

Due to the dependence of the single-colour decay factor on optical depth, it is optimal to have two-colour stationary light in an XPM scheme. This means that only one of the light fields will generate a XPS. The use of single-colour dispersion-free stationary light is not considered. The success probability is improved due to both signal fields producing a shift with a corresponding reduction in the decay of the signal state. However the single-colour stationary light operates at reduced effective optical depth due to interference between atomic transitions, which likely cancels out any advantage.

13.2 Performance of a phase gate based on Raman stationary light

The calculation of Sec. 13.1 is not complete for use of the nonlinearity in the quantum regime. It ignores the probabilistic nature of the decay of the two states. It suggests that with an optical depth of $d \sim 100$ and sufficient focussing, which is a difficult but not impossible combination, a phase shift $\phi_p = \pi$ should be possible. This magnitude of phase shift has not yet been demonstrated using stationary light, and a more thorough examination of the interaction will show that a much larger optical depth is necessary for such a demonstration.

13.2.1 Success probability for a simplified gate

An important additional interaction between the stationary light field and the target state must be included. The signal field is close to resonance with a populated transition from the target state. Even if there is just one atom in that state, there is a chance of scattering on that transition and destroying both states. The proportionality of Δ_S between this scattering rate and the AC-Stark shift indicates that detuning further would be advantageous for preserving the states. However, this can reduce the total cross-phase shift by increasing the chance of the signal state decaying. There is then a detuning which maximises the cross-phase shift, that increases with total optical depth.

In addition to the chance of scattering off the target state, the signal field also experiences dispersion while crossing the target state near to resonance. The phase shift of the light will then cause an evolution of the phase of the stationary coherence. These are proportional, as the phase shift of the signal field at $\Delta_S \gg \Gamma$ is proportional to $Re[\chi]$ while the AC-Stark shift is proportional to $Re[\chi]|\hat{\mathcal{E}}|^2$. This is not a coincidence; the dispersion results from coherent emission by the excited target state atoms, while the AC-Stark shift results from the atoms spending time partially in the excited state. In fact, the corresponding phase shift of the signal state is shown below to be proportional to the AC-Stark shift and so the cross-phase shift of the target state.

The relationship between dispersion and AC-Stark shift is also suggested by quantum mechanics: It is not possible to induce with one state a phase shift on another without some corresponding back-action, as this would allow a measurement with accuracy exceeding the Heisenberg limit. There is no need for the phase shifts of each state to be exactly equal. However, in the geometry considered and where the interaction can be made practically lossless, this constraint is applied to avoid allowing unphysical measurements.

The phase shift of the signal state is also necessary for implementing the desired gate, in which a phase shift of each state occurs only in the case that both states are present. We calculate the dependence of the evolution of each state on the basic experimental parameters.

To build the gate, the states are arranged in the memory as in Fig. 13.1. The signal state is stored over the maximal amount of optical depth, in the shape of a stationary spinwave over regions I and III. The target state will be stored over a small amount of optical depth in the central region II, where the stationary light will have maximum intensity. We consider the case where each state has a single quantum, and set the regions of the signal state $S_1 = -S_3 = 1$, and the target state $T_2 = 1$. Note that setting these quanta equal to 1 is a slightly different renormalisation of the Maxwell-Bloch equations to that used in

previous calculations.

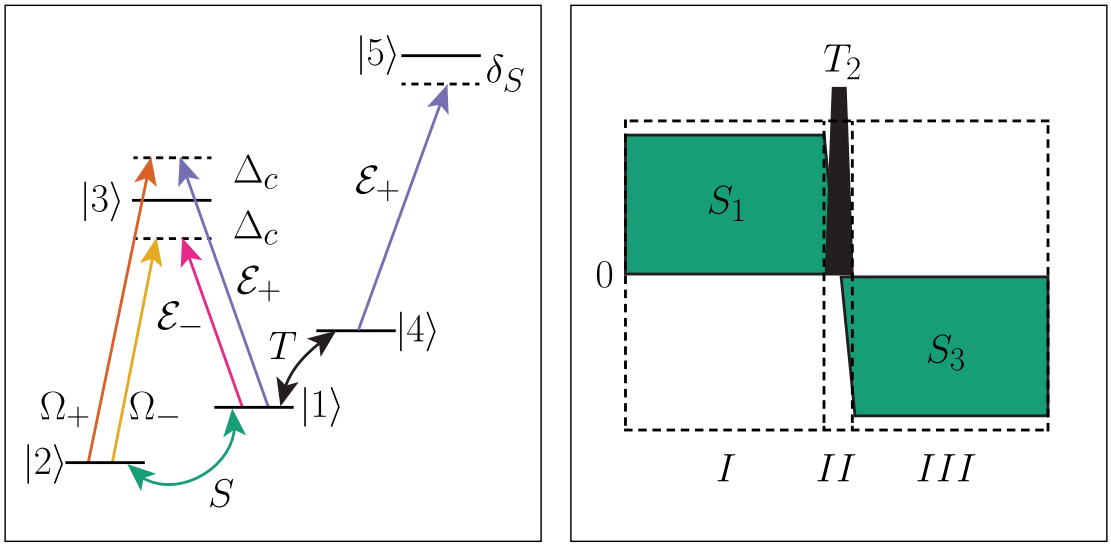

Fig. 13.1. Left, simplified level scheme with stationary light from spinwave S on levels $|1-3\rangle$ and one stationary field modulating level $|4\rangle$ on which T is stored. Right, the two states are stored over different regions of the ensemble.

Setting the signal field so that $\delta_S = \Gamma/2$, the phase shift ϕ of the field crossing the target state (which consists of a single atom in state $|4\rangle$) is approximately

$$\begin{aligned}\phi &= \frac{\Gamma\sigma_{at}}{4A\delta_S} \\ &= \frac{\sigma_{at}}{2A}.\end{aligned}$$

The AC-Stark shift of the atom is

$$\begin{aligned}\Delta_{AC} &= -\frac{\Omega_{sig}^2}{4\delta_S} \\ &= -\frac{\sigma_{at}}{2A}\left|\hat{\mathcal{E}}_+\right|^2, \end{aligned} \tag{13.10}$$

where $\hat{\mathcal{E}}_+ = \frac{\sqrt{d\Gamma}}{2}\frac{\Omega_c}{\Delta_c}S_1$, with the factor of $\sqrt{\Gamma}$ resulting from the renormalisation required to set a single quanta to $|S|^2 = 1$ instead of $|S|^2 = \Gamma$. This gives the AC-Stark shift as

$$\begin{aligned}\Delta_{AC} &= \frac{\sigma_{at}}{2A}\frac{d\Gamma\Omega_c^2}{4\Delta_c^2}|S_1|^2 \\ &= \frac{\Omega' d\phi}{2}|S_1|^2, \end{aligned} \tag{13.11}$$

where $\Omega' = \frac{\Gamma\Omega_c^2}{2\Delta_c^2}$ with d the resonant optical depth of the ensemble on the transition $|1\rangle \rightarrow |3\rangle$.

Meanwhile, S will evolve according to a slightly modified stationary light equation. The evolution of S_1 is proportional to the integral of the spinwave (or sum in this case). Since the phase of some of the forward propagating optical field in S_3 has been shifted,

the S_1 portion of the sum is simply multiplied by the appropriate phase.

$$\partial_t S_1 = -\frac{d\Omega_c^2\Gamma}{2\Delta_c^2}\left(S_1 + S_3\right) \tag{13.12}$$

$$\partial_t S_3 = -\frac{d\Omega_c^2\Gamma}{2\Delta_c^2}\left(e^{i\phi}S_1 + S_3\right) \tag{13.13}$$

Assuming that the optical depth of region II is much smaller than the other regions, the evolution of the signal state coherence in this region can be ignored. The equations of motion become

$$\partial_t S_1 = -\Omega' d(S_1 + S_3) - \Omega' S_1 \tag{13.14}$$

$$\partial_t T_2 = -i\frac{\Omega' d\phi}{2}|S_1|^2 T_2 \tag{13.15}$$

$$\partial_t S_3 = -\Omega' d(e^{i\phi}S_1 + S3) - \Omega' S_3 \tag{13.16}$$

These equations of motion, in the limit of high optical depth, describe a high fidelity cross-phase gate on the states T and S. The mechanism by which the phase of S is shifted is hidden in the equations: For small ϕ, a phase shifted component of the signal state $-i\Omega' d\phi S_1$ is formed in region III. Due to the rapid self-stabilising evolution, this is spread out across region I. This component then grows over both regions at $\mp i\Omega' d\phi S_1/2$, the same rate at which the phase of the target state changes.

Due to the phase shift of the signal field, a smaller real component of opposite sign to S_3 is produced. This decreases the magnitude of S_3 and in turn S_1. This is the second loss mechanism apart from control field scattering, and can be considered as the probability that the signal photon escapes rather than being absorbed into the spinwave. The loss rate is $\Omega'\phi^2 d/4$, while the phase shift is proportional to $\Omega'\phi d/2$. The loss over a pi phase shift is $\pi\phi/2$.

Losses will affect the performance of the gate, and these can be determined analytically with the simplified dynamics. As long as the losses are small, the magnitudes of the signal and target state indicate the success rates of the gate. However, when using the semiclassical equations the reduction in amplitude of the states affects the amount of light produced and thus the phase shifts. This is not realistic, as when single photons are used the loss of any of the states would count as a failure, and conversely the phase accumulated in a successful gate would not be affected by *any* loss process. In this case it is appropriate to simply add each loss rate (as the mechanisms are mutually exclusive) to determine the final success probability.

The above can be arranged to find the losses for the various processes over a π phase shift of both states. The more accurate AC-Stark shift proportional to $\Gamma\delta_S/(4\delta_S^2 + \Gamma^2)$ is used, and the cross section $\sigma_{at}/A = \sigma'$.

The effective losses for the different mechanisms are written as rates scaled to the dimensionless time $t\Delta_{AC} = \pi$, with the loss due to control field scattering

$$\gamma_\Omega = \frac{2\pi}{d}\frac{4\delta_S^2 + \Gamma^2}{\Gamma\delta_S}, \tag{13.17}$$

the escape probability for the signal photon

$$\gamma_S = \frac{\pi\Gamma\delta_S\sigma'}{4\delta_S^2 + \Gamma^2}, \tag{13.18}$$

and the scattering probability of the signal field from the target state

$$\gamma_T = \frac{\pi\Gamma}{\delta_S}. \tag{13.19}$$

These rates are multiplied to give the success probability for a π phase shift gate,

$$P_{Success} = \exp(-\gamma_\Omega - \gamma_S - \gamma_T). \tag{13.20}$$

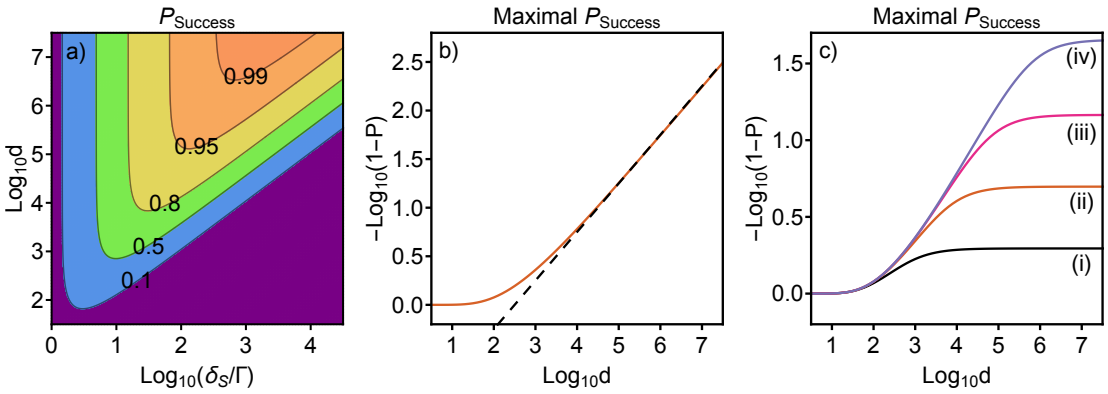

Fig. 13.2. Controlled phase gate success probabilities: a) As a function of optical depth and detuning with $\sigma' = 0.5$. b) As a function of optical depth (orange) with dashed line the asymptotic gradient of 0.5. c) As a function of optical depth where the focussing parameter σ' is constrained by assuming ensembles with various transverse extents: (i) $\sigma' = d/50$; (ii) $\sigma' = d/500$; (iii) $\sigma' = d/5000$; (iv) $\sigma' = d/50000$.

The success rate as a function of optical depth d and detuning δ_S are plotted in Fig. 13.2, with the assumption that the size of the interaction volume can be on the order of a wavelength ($\sigma' = 0.5$). The plotted line is calculated from Eq. (13.20), which also asymptotes to $P_{Success} = \exp\left(-18/\sqrt{d}\right)$. In Fig. 13.2 c), the interaction volume is constrained by setting $\sigma' \propto d$. This is equivalent to making assumptions about the size and density of ensembles that can be produced. The gradient of maximal $P_{Success}$ as a function of optical depth asymptotes to $1/2$ for the small interaction volume. This means that the number of successful gates that can be performed scales only with $\sqrt{d}$, even though the magnitude of the achievable cross-phase shift scales with d.

Enhancement of a free-space interaction is impeded by the rapid divergence of tightly focussed beams. Raman stationary light allows use of optical depth outside the Rayleigh length, but there can be experimental limitations to the amount of optical depth available. The success rates in Fig. 13.2 c) are calculated with the cross-sectional area of the interaction region inversely proportional to optical depth. A selection of different proportionality constants are used. The constraint roughly corresponds to a situation where optical depth can only be increased by increasing the length of the ensemble with a fixed transverse distribution, or alternatively by further compressing a fixed number of atoms. This indicates that the success probability is not purely dependent on the optical depth, but

that the geometry is also important. Note that these calculations are still one-dimensional and do not include the effects on phase-matching of the three-dimensional beam.

The escape probability of the signal photon can be identified with the causal limitations discussed in [196, 197]. A large phase shift of the signal photon crossing the target state allows the photon to escape, which is analogous to the problem in EIT-enhanced XPM where the photon is pushed outside the EIT bandwidth and absorbed. This also points to the solution of the problem in [198], where it is conjectured that many photons must be sent to realise a high fidelity phase shift. Instead, here, as in a cavity, the same photon travels across the interaction region many times.

The comparison to a cavity is an interesting one, as increasing the finesse of a cavity involves decreasing its bandwidth (after the cavity has been shrunk to the smallest possible size) and so lowering the rate of operations possible in cavity-enhanced XPM. The question of whether this bandwidth limitation could be overcome in XPM with Raman-SL could be answered by going beyond the steady-state equations used above. This will be left for future work.

13.2.2 Fidelity of a successful gate

As mentioned in Sec. 13.2.1, the semi-classical equations of motion do not give entirely accurate results for quantum states as they model the decay of the states throughout the interaction. A gate is only considered successful if the states are not lost during the interaction, so no decay should occur. The semi-classical equations are however used in the far-detuned regime to investigate the effect of the longitudinal distributions of the signal and target states on the fidelity of the gate.

An extra interaction is neglected in the success rate calculations: the effect of the dispersion of the signal field on the *distribution* of the signal state. The above calculations limited the signal state spinwave to a symmetric, two point distribution separate from the target state. For a target state with longitudinal extent, the spatially varying phase of the signal field can produce a spatially varying evolution of the signal state. Changing the spatial distribution of the signal state (conditional on the presence of a target photon) would necessarily change the spectral distribution of the output signal photon, lowering the gate fidelity and producing spectral entanglement.

One dimensional simulations were run, with the implicit assumption of a uniform interaction in the transverse dimensions. All loss mechanisms that could be discarded were. Including the signal field dispersion left the possibility of signal photon escape, so the parameters were chosen to minimise this probability. This parameter selection did not affect the shape-dependent interaction.

The simplified equations of motion solved in the simulations were:

$$\partial_t S = i\Omega \left(\hat{\mathcal{E}}_+ + \hat{\mathcal{E}}_-\right) \tag{13.21}$$

$$\partial_t T = -i\theta \left|\hat{\mathcal{E}}_+\right|^2 T \tag{13.22}$$

$$\partial_z \hat{\mathcal{E}}_+ = i\left(\Omega S + \theta \left|T\right|^2 \hat{\mathcal{E}}_+\right) \tag{13.23}$$

$$\partial_z \hat{\mathcal{E}}_- = -i\Omega S \tag{13.24}$$

These preserve the requirement for equal phase shifts on the two states (assuming

$\int |S|^2 \, \mathrm{d}z = \int |T|^2 \, \mathrm{d}z = 1$) while simplifying the calculations. As in the success probability calculations, for small θ the phase shift in the signal state is introduced by the component of the forward-travelling signal field that is shifted while crossing the target state. For shape-preserving evolution, the magnitude of the spinwave at any point should approximately match the magnitude of this shifted signal field, giving

$$S(z) \propto \left(\hat{\mathcal{E}}_+(z) + \hat{\mathcal{E}}_-(z)\right) \rightarrow S(z) \propto \int_0^z \left|T(z')\right|^2 \mathrm{d}z' \tag{13.25}$$

This is only an approximate solution, as it ignores the probe light emitted in the interaction region. The phase shift of this light lags behind the rest and the spinwave shape in the region is altered. This error cannot be corrected for with a stationary spinwave, as setting the phase of the spinwave proportional to the integral $\int |T(z')|^2 \, \mathrm{d}z'$ gives a non-stationary spinwave. Instead, the error can be minimised by minimising the proportion of the ensemble over which the target state is stored. This also improves the uniformity of the phase shift of the target state by minimising the variation of the intensity of the signal field over the target state.

The two distributions shown in Fig. 13.3 have target states that are Gaussians of different widths. The shape preserving evolution works well even at the larger width of the target state, but there is measurable variation of the shape in the interaction region. The phase variation across the target state is more noticeable for the wider target state, and significantly smaller for the narrower target state.

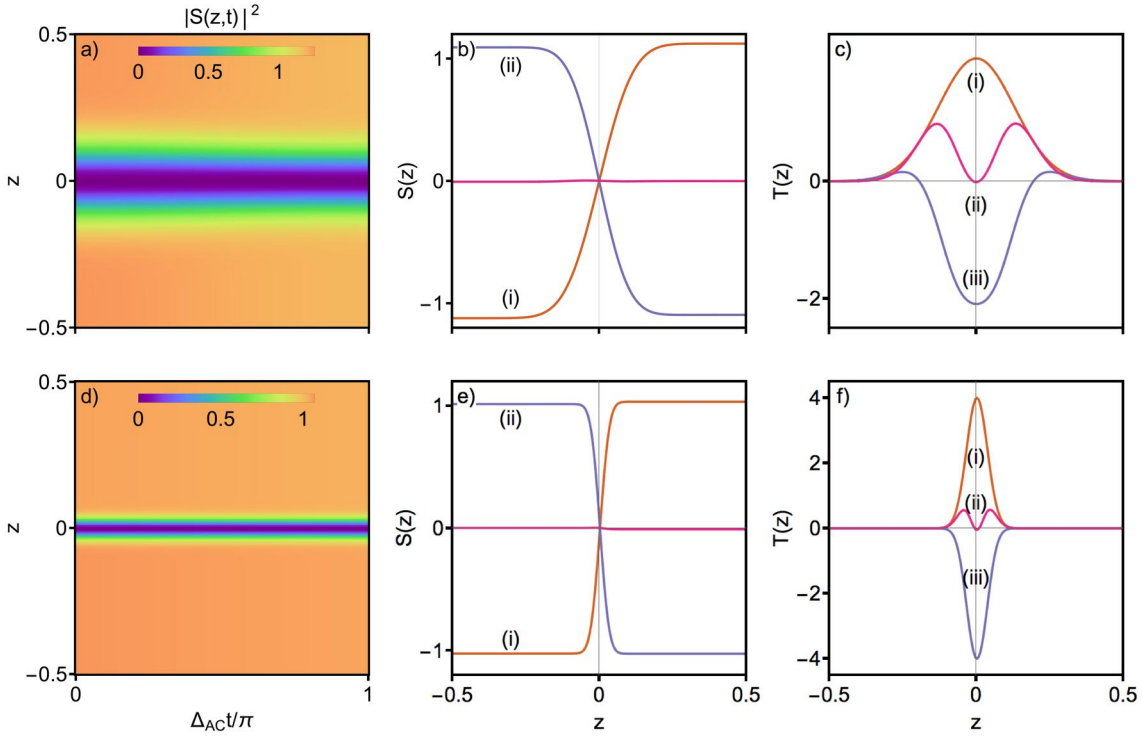

Fig. 13.3. Semiclassical simulation results for two different target state widths. a) and d) The spinwave distributions over time. b)(i),e)(i) The initial spinwaves $S(z,0)$, and (ii) the final spinwaves $S(z,1)$. c)(i),f)(i) initial spinwaves T(z,0), final spinwave (ii) imaginary component $Im(T(z,1))$, and (iii) real component $Re(T(z,1))$.

A fidelity can be approximated for a successful gate by calculating the overlap of the states with the desired output states of the gate. Since this is still a semiclassical calculation, the two-particle operator of Eq. (12.4) cannot be used. Instead, the overlaps of each individual state with their desired output state are multiplied together. This overestimates the fidelity as the spectral entanglement cannot be accounted for. Fig. 13.4 plots the estimated fidelity $\mathcal{F}$ versus the width of the target state in a regime where losses do not affect the fidelity.

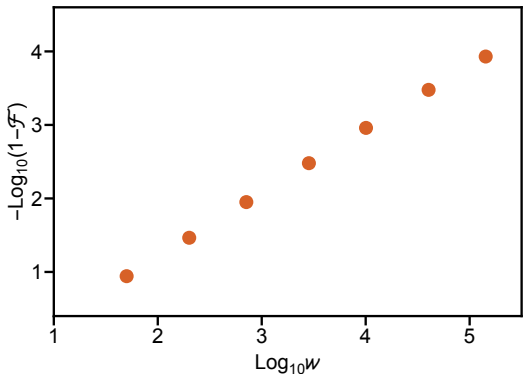

Fig. 13.4. Approximate fidelity from overlap integrals for a successful π phase gate. The target state is a Gaussian with $T(z) = w\sqrt{2/\pi}\exp\left(-(wz)^2\right)$.

There is then a limit to fidelity in this technique depending on the proportion of the memory used to store the target state. However, minimising this portion will decrease the storage and retrieval efficiency of the target state, reducing the gate success rate.

An alternative scheme would involve piecewise illumination of the ensemble by control fields. With no overlap between the signal and target states, there would be no corresponding loss of fidelity. The intensity of the signal field would not vary over the region containing the target state, and the phase of the signal field would not vary over the region containing the signal state. Calculating the final state via the two-particle operator (12.4) would also not result in unwanted entanglement due to the constant distribution of the signal field in the region of the target state.

We can compare this to the entanglement problem for EIT based schemes. The EIT interaction is sensitive to any phase shift of the optical component, and so the presence of the target state modifies the propagation of the signal state locally. The relationship between position and frequency means entanglement is generated between frequency modes. For Raman SL, the evolution of the signal spinwave in the region of the target state depends on the position of the target state, because it is at that position that the phase shift occurs. This would then generate the same sort of entanglement. However, because the evolution of the signal spinwave can be switched off in the target region (or just minimised by reducing the optical depth of the target region), this entanglement can be entirely avoided.

In the future, adapting the two-particle operator evolution approach to a Raman stationary light cross-phase gate would allow more accurate calculations of gate fidelities. Looking further ahead, this theory could be extended to test the effect on the fidelities and success rates of having two- or three-dimensional distributions in the states, in preparation for a demonstration in the setup used in the stationary light experiments. The signal field mode will experience distortion due to phase variation in the transverse dimension

in signal spinwave, caused by the transverse distribution of the target state. In addition to the limitations of three-dimensional stationary light, this will place very stringent requirements on the shapes of the states in the transverse dimension.

13.3 Optical switch based on Raman stationary light

Another protocol is possible with a very similar setup. Two extra ingredients are necessary: The ability to create a large conditional dispersion - a phase shift of π on the signal fields over a single pass, and equal rather than opposite detunings for the signal fields. The large dispersion of the signal field changes the stationary light condition. For example, a signal spinwave written with uniform phase across the memory would be trapped rather than escape when the control fields are turned on. The constructive interference of the counter-propagating signal fields into the spinwave coupling is converted to destructive interference as the phase of the two fields is shifted by π while crossing the target state. One problem with this scheme is that for single photon states, a dispersive phase shift of π in a single crossing of the target state is only possible where the signal field is on resonance with the target transition. This gives a very high probability of incoherent scattering. However, state-dependent destruction of photons can also be a desirable effect.

This scheme could also be performed based on the conditional storage of the signal state in a TRACE-like scheme. The interference of the counter-propagating fields would be converted from constructive to destructive by the dispersion, preventing storage of the input light.

13.4 Swap gate based on Raman stationary light

If the signal state does not occupy the whole ensemble, the signal state can be redistributed into the empty areas. The phase shifted component of the field is not cancelled by any other field anywhere in the ensemble, and so a spinwave will be generated in any empty areas that are illuminated by the control fields. This newly generated spinwave then affects the evolution in the initial region, producing a more complicated overall evolution of the spinwave. The interaction between the separate regions of spinwave is purely linear, suggesting that the linear parts of quantum computing gates could also be implemented within the memory.

The interaction is demonstrated in Fig. 13.5, where empty regions exist outside the region in which the signal state is stored. For the time in which the initial region would normally accumulate a π phase shift, the spinwave migrates entirely to the unoccupied regions. If the evolution is allowed to continue, the spinwave migrates back to its initial position. If these regions encode two separate qubits, this implements a control-swap, or Fredkin gate.

Ignoring the scattering decays, the evolution can again be solved along the lines of equations (13.16). For simplicity we now assume a phase shift is applied to probe fields travelling in *both* directions across the target state. The encoding requires stationary (in the absence of a target state) spinwaves. For example, $A_1(0) = -A_2(0) = \psi_A$, and $B_1(0) = -B_2(0) = \psi_B$.

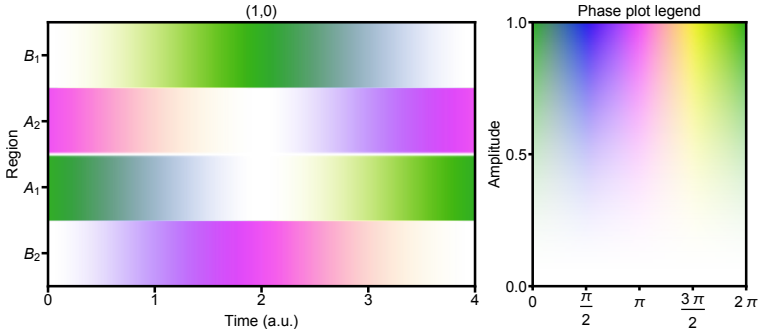

Fig. 13.5. Illustration of a control-swap gate. A target state is present between A_1 and A_2. Over an interval of time 2, the memory performs a swap between regions A and B. Over an interval of time 1, $\sqrt{SWAP}$ is applied. The phase of the spinwave is indicated by colour, and the magnitude is indicated by opacity.

Applying the control fields gives the evolution

$$\frac{d}{dt}A_1(t) = -\Omega' d\left(A_1(t) + B_2(t) + \exp(i\phi)\left(A_2(t) + B_1(t)\right)\right) \tag{13.26}$$

And so on for the remaining regions of spinwave. Since the evolution is symmetric due to the symmetric phase shift, set $A_1(t) = -A_2(t) = \psi_A(t)$ and $B_1(t) = -B_2(t) = \psi_B(t)$. These ψ represent the states of the qubits, but we are leaving out the ket representation since it is unnecessary for this calculation. Solving the equations of motion gives

$$\psi_A(t) = \frac{1}{2}\left(\psi_A(0) + \psi_B(0) + \exp\big(2t\Omega' d(\exp(i\phi) - 1)\big)\left(\psi_A(0) - \psi_B(0)\right)\right)\Big) \tag{13.27}$$

$$\psi_B(t) = \frac{1}{2}\left(\psi_A(0) + \psi_B(0) + \exp\big(2t\Omega' d(\exp(i\phi) - 1)\big)\left(\psi_B(0) - \psi_A(0)\right)\right)\Big) \tag{13.28}$$

For $\phi \ll 1$, the term $\exp(i\phi) - 1 = i\phi - \phi^2/2$. The real-valued ϕ^2 term is the decay due to signal photon escape, which is discussed above. For sufficiently small ϕ the ϕ^2 term can be ignored (the loss is proportional to $1 - \exp(-\pi\phi)$ over a complete swap gate). Where $\psi_A(0) = \psi_B(0)$, there is no evolution of the states, as no light crosses the control qubit. The swap gate is completed at $t = 1/(2\Omega' d\phi)$, where $\psi_A(t) = \psi_B(0)$ and $\psi_B(t) = \psi_A(0)$. This replicates the interaction suggested in [212]. The evolutions for states $\psi_A(0 = \psi_B(0)$, $\psi_A(0) = -\psi_B(0)$ and $\psi_A(0) = i\psi_B(0)$ are shown in Fig. 13.6.

Where the two regions encode the two parts of a single qubit, the interaction can be used to implement any controlled unitary rotation, provided that the interaction in the regions can be controlled individually (for example, by illuminating each pair of regions with control field for different durations or at different intensities). Unconditional gates can also be implemented by simulating the presence of a target state. This is done by slightly shifting the phases of the control fields between regions. As shown in Fig. 13.2, a large optical depth is required to implement a reliable gate. Therefore, unconditional gates should instead be implemented by retrieving the states and interfering them with the desired phase, which can be done in a classical interferometer or by one of the many appropriate techniques discussed in this work.

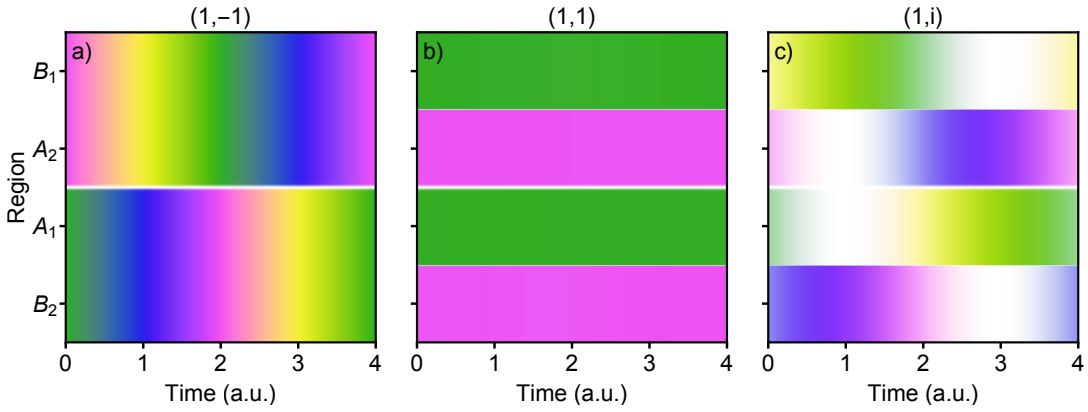

Fig. 13.6. Control-swap gate for various inputs. In (a) the two states are π out of phase. The swap gate resembles a π phase gate as it flips the phase of each input. In (b) the inputs are identical and so the swap gate has no effect. Physically, there is no nonlinear interaction as no light crosses the target state. In c) the inputs are $\pi/2$ out of phase, emphasising the $\sqrt{SWAP}$ gate which is performed by time 1 and adds the two original states $\pm\pi/2$ out of phase.

13.5 Conclusion

In this chapter we have touched upon basic atom-light interactions such as AC-Stark shift and dispersion, and how they relate to XPM. We have shown how the integrated intensity of light generated in Raman stationary light is dependent on the optical depth of the ensemble. We have also shown that the dispersion of the stationary light field crossing an atom in the target state limits the available cross-phase shift, reducing the dependence of gate success to proportional to $\sqrt{d}$. We have shown how a control-swap gate may be implemented entirely inside an atomic ensemble.

Chapter 14

Conclusion and outlook

> Patience is bitter but its fruit is mad sweet,
> Like a swole grape.
>
> Jean-Jacques Rousseau,
> as quoted in *Welcome to Nightvale*

In the first part of this thesis we gave examples of how efficient and versatile a gradient echo memory can be. We demonstrated and characterised an optical quantum memory in cold rubidium atoms. This was the first memory shown to outperform the theoretical limit for storage of quantum states in an optical fibre loop.

We also used the gradient echo memory to implement a pulsed optical resonator in a warm rubidium vapour. This resonator behaved in a predictable way based on a few experimental parameters which were rapidly tunable.

The case of a one dimensional gradient echo memory was investigated. We studied two extensions that use multiple control fields at different angles to increase the flexibility of the memory. The first scheme is capable of retrieving stored states without necessarily having to switch the gradient and rapidly reordering the stored states upon retrieval. The second scheme recreates a simulated gradient by combining a number of static control fields. This results in a memory that sits conceptually somewhere between a gradient echo memory and a Raman memory.

The gradient echo memory is a versatile and effective protocol for storage of optical quantum states. In the near future, work in GEM will involve further improving the efficiency and storing single photons. In the longer term we hope that something resembling GEM will be used as part of the quantum internet, which would require the development of repeater protocols, as well as considerable increases in efficiency, scalability, and compatibility with telecommunications wavelengths.

In the second part of this thesis we proposed a new form of stationary light, Raman stationary light, and demonstrated it in cold atoms. This form of stationary light is significantly different from previous forms as it relies on a spatially extended interference effect. This means that the stationary light is dependent on the phase and spatial distribution of the spinwave within the memory. The behaviour of the spinwave also allowed a single-mode memory to be implemented. We investigated how the efficiency of this memory compared to other single-mode memories. We also investigated the temporal mode selectivity of single-mode memories more generally. We then demonstrated the memory in cold rubidium atoms, investigating some of the additional complications that arise due to dispersion and the Gaussian beam shapes of the stored optical fields.

In the third part of this thesis we investigated the enhancement of cross-phase modu-

lation by using Raman stationary light. We found that the causal limit to many ensemble enhanced interactions also applies to Raman stationary light, making it a considerable challenge to generate π cross-phase shifts between single photons in this scheme. Additionally, the D1 line of rubidium-87 which provides the level structure for many of our experiments is not practical for performing a true cross-phase experiment. However, the fact that the stationary light can be spatially separated from the spinwave should allow this scheme to avoid additional entanglement-related noise inherent to EIT-based stationary light. It will be useful to continue exploring the idea, focussing on experimental implementations and with a mind to designing scalable systems. From a theoretical perspective, a more thorough quantum treatment of the cross-phase modulation should be possible. The transient behaviour of the nonlinearity as well as stationary light could be considered in order to answer the question of whether bandwidth limitations for cross-phase modulation in cavity-ensemble systems also apply to the Raman stationary light enhanced modulation.

We have shown that atomic ensembles are useful as an interface for storing and processing quantum states of light. This will continue to be true for some time. As we have seen there is a huge amount of activity in developing these systems for their expected applications, and in discovering new applications.

Appendices

Appendix A

XMDS2 simulation

Here we include an example XMDS2 script. More examples of scripts based on those used to generate simulations for this thesis, along with visualisation tools, can be found at https://github.com/jlleverett/AtomOpticMemorySims.git

A.1 Gradient echo memory simulation

By including the complex detuning $\tilde{\Delta}$, a two-level simulation closely approximates a three-level simulation, so long as the adiabatic approximation is valid.

```
<?xml version="1.0" encoding="UTF-8"?>
<simulation xmds-version="2">
<name>quasi3levelstorage</name>
<author>Jesse</author>
<description>
2-level simulation of GEM with adiabatic approx
</description>

<features>
<benchmark /><bing /><fftw plan="exhaustive"/><auto_vectorise />
<globals><![CDATA[
//Physical constants
const real pi = M_PI;

//Experimental parameters: control field Rabi freq = om,
//excited state linewidth = gama, one-photon detuning = deltaq
//bandwidth = grad/2pi, optical depth = d.
 double gama = pi*5.75;
 double deltaq = 2*pi*440;
 complex delta = (gama*gama+deltaq*deltaq)/(deltaq + i*gama);
 double grad = 2*pi*0.1;
 double om = 2*pi*10;
 double d = 600;
 complex omf; double eta;

// Defines a gaussian with width w.
 double gaussian( double x, double w ) {
 return exp( -2*pow(x,2)/(pow(w,2)) ); }
```

```
// Control field switches as functions of time.
 double cswitch(double time){
 return (time < 30) ? 1.0 : (time < 40) ? 0.0 : 1.0;}

//Similarly for the gradient
 double gswitch(double time){
 return (time < 30) ? 1.0 : (time < 40) ? 0.0 : -1.0;}
]]></globals></features>

<geometry>
<propagation_dimension> t </propagation_dimension>
<transverse_dimensions>
<dimension name="z" lattice="100"  domain="(-0.5,0.5)" />
</transverse_dimensions></geometry>

<vector name="main" initial_space="z" type="complex">
<components>S</components>
<initialisation>
<![CDATA[S = 0;]]>
</initialisation></vector>

<vector name="cross" initial_space="z" type="complex">
<components>E</components></vector>

<sequence>
<integrate algorithm="ARK45"
interval="80" steps="400" tolerance="1.0e-9">
<samples>400</samples>
<operators>
<operator kind="functions">
<![CDATA[
// Defines the control fields and gradients.
 omf = om * cswitch(t);
 eta = grad * gswitch(t);
]]>
</operator>

<operator kind="cross_propagation"
algorithm="RK4"  propagation_dimension="z">
<integration_vectors>cross</integration_vectors>
<dependencies>main</dependencies>
<boundary_condition kind="left">
<![CDATA[
//Defines probe input profile.
 E=gaussian(t-16,12);
]]>
</boundary_condition>
```

```
<![CDATA[
//spatial DE for probe field
dE_dz = i*sqrt(d)*omf/delta*S;
]]>
</operator>

<integration_vectors>main</integration_vectors>
<dependencies>cross</dependencies>
<![CDATA[
 dS_dt=i*eta*z*S+i*omf*gama/delta*E;
]]>
</operators></integrate></sequence>

<output format="hdf5" filename="quasi3levelstorage.xsil">
<group>
<sampling basis="z" initial_sample="yes">
<moments>SR SI ER EI</moments>
<dependencies>main cross</dependencies>
<![CDATA[
 _SAMPLE_COMPLEX(S);
 _SAMPLE_COMPLEX(E);
]]>
</sampling></group></output>
</simulation>
```

Bibliography

[1] Lambropoulos, and D. Petrosyan. *Fundamentals of Quantum Optics and Quantum Information*. Springer-Verlag, 2007.

[2] D. Walls, and G. Milburn. *Quantum Optics*. Springer-Verlag, 1994.

[3] U. Leonhardt, and H. Paul. Measuring the quantum state of light. *Progress in Quantum Electronics*, 19(2):89–130, 1 1995.

[4] M. Orszag. *Quantum Optics*. Springer International Publishing, Cham, 2 edition, 2016.

[5] U. Leonhardt, and H. Paul. Phase measurement and Q function. *Physical Review A*, 47(4):2460–2463, 1993.

[6] D. A. Steck. *Quantum and Atom Optics*. 2012.

[7] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen. Photon storage in Λ-type optically dense atomic media. II. Free-space model. *Physical Review A*, 76(3):033805, 9 2007.

[8] M. Fleischhauer, A. Imamoglu, and J. P. Marangos. Electromagnetically induced transparency: Optics in coherent media. *Reviews of Modern Physics*, 77(2):633–673, 7 2005.

[9] M. Fleischhauer, and M. Lukin. Dark-state polaritons in electromagnetically induced transparency. *Physical review letters*, 84(22):5094–7, 5 2000.

[10] J. Nunn et al. Enhancing multiphoton rates with quantum memories. *Physical Review Letters*, 110(13):1–5, 2013.

[11] A. Gorshkov, A. André, M. Fleischhauer, A. Sørensen, and M. Lukin. Universal Approach to Optimal Photon Storage in Atomic Media. *Physical Review Letters*, 98(12):123601, 3 2007.

[12] F. Kaneda et al. Time-multiplexed heralded single-photon source. *Optica*, 2(12):1010, 12 2015.

[13] J.-i. Yoshikawa, K. Makino, S. Kurata, P. van Loock, and A. Furusawa. Creation, Storage, and On-Demand Release of Optical Quantum States with a Negative Wigner Function. *Physical Review X*, 3(4):041028, 12 2013.

[14] N. Lundblad, M. Schlosser, and J. V. Porto. Experimental observation of magic-wavelength behavior of 87Rb atoms in an optical lattice. *Physical Review A*, 81(3):031611, 3 2010.

[15] S.-J. Yang, X.-J. Wang, X.-H. Bao, and J.-W. Pan. An efficient quantum light-matter interface with sub-second lifetime. *Nature Photonics*, 10(6):381–384, 6 2016.

[16] M. Zhong et al. Optically addressable nuclear spins in a solid with a six-hour coherence time. *Nature*, 517(7533):177–180, 2015.

[17] M. Rančić, M. P. Hedges, R. L. Ahlefeldt, and M. J. Sellars. Coherence time of over a second in a telecom-compatible quantum memory storage material. *Nature Physics*, 14(1):50–54, 9 2017.

[18] D. G. England, P. J. Bustard, J. Nunn, R. Lausten, and B. J. Sussman. From photons to phonons and back: A THz optical memory in diamond. *Physical Review Letters*, 111(24):1–5, 2013.

[19] E. Poem et al. Broadband noise-free optical quantum memory with neutral nitrogen-vacancy centers in diamond. *Physical Review B - Condensed Matter and Materials Physics*, 91(20):1–10, 2015.

[20] L. Hau, S. Harris, Z. Dutton, and C. Behroozi. Light speed reduction to 17 metres per second in an ultracold atomic gas. *Nature*, 397(February):594–598, 1999.

[21] D. F. Phillips, A. Fleischhauer, A. Mair, R. L. Walsworth, and M. D. Lukin. Storage of Light in Atomic Vapor. *Physical Review Letters*, 86(5):783–786, 1 2001.

[22] M. Fleischhauer, and M. Lukin. Quantum memory for photons: Dark-state polaritons. *Physical Review A*, 65(2):022314, 1 2002.

[23] J. F. Marangos. Electromagnetically induced transparency. *Journal of Modern Optics*, 45(3):471–503, 1998.

[24] T. Chanelière et al. Storage and retrieval of single photons transmitted between remote quantum memories. *Nature*, 438(7069):833–6, 12 2005.

[25] M. D. Eisaman et al. Electromagnetically induced transparency with tunable single-photon pulses. *Nature*, 438(7069):837–841, 12 2005.

[26] S. Zhou et al. Optimal storage and retrieval of single-photon waveforms. *Optics express*, 20(22):24124–31, 2012.

[27] P. M. Anisimov, J. P. Dowling, and B. C. Sanders. Objectively Discerning Autler-Townes Splitting from Electromagnetically Induced Transparency. *Physical Review Letters*, 107(16):163604, 10 2011.

[28] L. Giner et al. Experimental investigation of the transition between Autler-Townes splitting and electromagnetically-induced-transparency models. *Physical Review A*, 87(1):013823, 1 2013.

[29] Y. H. Chen et al. Coherent optical memory with high storage efficiency and large fractional delay. *Physical Review Letters*, 110(8):1–5, 2013.

[30] J. Geng et al. Electromagnetically induced transparency and four-wave mixing in a cold atomic ensemble with large optical depth. *New Journal of Physics*, 16(11):113053, 11 2014.

[31] G. Heinze, C. Hubrich, and T. Halfmann. Stopped Light and Image Storage by Electromagnetically Induced Transparency up to the Regime of One Minute. *Physical Review Letters*, 111(3):033601, 7 2013.

[32] D. Akamatsu, K. Akiba, and M. Kozuma. Electromagnetically induced transparency with squeezed vacuum. *Physical Review Letters*, 92(20):1–4, 2004.

[33] D. Akamatsu et al. Ultraslow propagation of squeezed vacuum pulses with electromagnetically induced transparency. *Physical Review Letters*, 99(15):1–4, 2007.

[34] G. Hètet et al. Delay of Squeezing and Entanglement using Electromagnetically Induced Transparency in a Vapour Cell. *Optics Express*, 16(10):7369–7381, 2008.

[35] K. Honda et al. Storage and retrieval of a squeezed vacuum. *Physical Review Letters*, 100(9):4–7, 2008.

[36] K. S. Choi, H. Deng, J. Laurat, and H. J. Kimble. Mapping photonic entanglement into and out of a quantum memory. *Nature*, 452(7183):67–71, 2008.

[37] P. Vernaz-Gris et al. High Performance Raman Memory with Spatio-Temporal Reversal (in review).

[38] Z. Xu et al. Long lifetime and high-fidelity quantum memory of photonic polarization qubit by lifting zeeman degeneracy. *Physical Review Letters*, 111(24):1–6, 2013.

[39] Y. O. Dudin, L. Li, and A. Kuzmich. Light storage on the time scale of a minute. *Physical Review A - Atomic, Molecular, and Optical Physics*, 87(3):1–4, 2013.

[40] B. Gouraud, D. Maxein, A. Nicolas, O. Morin, and J. Laurat. Demonstration of a Memory for Tightly Guided Light in an Optical Nanofiber. *Physical Review Letters*, 114(18):180503, 5 2015.

[41] B. Albrecht, C. Clausen, C. Sayrin, P. Schneeweiss, and A. Rauschenbeutel. Storage of light in a nanofiber-trapped atomic ensemble. In *2015 European Conference on Lasers and Electro-Optics - European Quantum Electronics Conference, (Optical Society of America, 2015)*, page EA_2_5, 2015.

[42] B. Gouraud. *Optical nanofibers interfacing cold atoms. A tool for quantum optics. Universite Pierre et Marie Curie, Paris.* PhD thesis, 2016.

[43] E. Hahn. Spin Echoes. *Physical review*, 297(1946), 1950.

[44] N. Kurnit, I. Abella, and S. Hartmann. Observation of a Photon Echo. *Physical Review Letters*, 13(19):567–569, 1964.

[45] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin. Multimode quantum memory based on atomic frequency combs. *Physical Review A*, 79(5):052329, 5 2009.

[46] M. Gündoğan, P. M. Ledingham, K. Kutluer, M. Mazzera, and H. de Riedmatten. Solid State Spin-Wave Quantum Memory for Time-Bin Qubits. *Physical Review Letters*, 114(23):230501, 6 2015.

[47] C. Laplane et al. Multiplexed on-demand storage of polarization qubits in a crystal. *New Journal of Physics*, 18(1), 2016.

[48] P. Jobez et al. Towards highly multimode optical quantum memory for quantum repeaters. *Physical Review A*, 93(3):1–9, 2016.

[49] M. Sabooni, Q. Li, S. Kröll, and L. Rippe. Efficient Quantum Memory Using a Weakly Absorbing Sample. *Physical Review Letters*, 110(13):133604, 2013.

[50] P. Jobez et al. Cavity-enhanced storage in an optical spin-wave memory. *New Journal of Physics*, 16(8):083005, 8 2014.

[51] T. Zhong et al. Nanophotonic rare-earth quantum memory with optically controlled retrieval. *Science*, 357(6358):1392–1395, 9 2017.

[52] M. Gündoğan, M. Mazzera, P. M. Ledingham, M. Cristiani, and H. de Riedmatten. Coherent storage of temporally multimode light using a spin-wave atomic frequency comb memory. *New Journal of Physics*, 15(4):045012, 4 2013.

[53] P. Jobez et al. Coherent Spin Control at the Quantum Level in an Ensemble-Based Optical Memory. *Physical Review Letters*, 114(23):1–5, 2015.

[54] E. Saglamyurek et al. Quantum storage of entangled telecom-wavelength photons in an erbium-doped optical fibre. *Nature Photonics*, 9(2):83–87, 2015.

[55] J. Jin et al. Telecom-Wavelength Atomic Quantum Memory in Optical Fiber for Heralded Polarization Qubits. *Physical Review Letters*, 115(14):1–5, 2015.

[56] A. Seri et al. Quantum Correlations between Single Telecom Photons and a Multimode On-Demand Solid-State Quantum Memory. *Physical Review X*, 7(2):021028, 5 2017.

[57] S. Moiseev, and S. Kröll. Complete Reconstruction of the Quantum State of a Single-Photon Wave Packet Absorbed by a Doppler-Broadened Transition. *Physical Review Letters*, 87(17):173601, 10 2001.

[58] B. Lauritzen et al. Telecommunication-wavelength solid-state memory at the single photon level. *Physical Review Letters*, 104(8):1–4, 2010.

[59] B. Lauritzen, J. Minář, H. De Riedmatten, M. Afzelius, and N. Gisin. Approaches for a quantum memory at telecommunication wavelengths. *Physical Review A - Atomic, Molecular, and Optical Physics*, 83(1):1–12, 2011.

[60] B. Julsgaard, and K. Mølmer. Fundamental limitations in spin-ensemble quantum memories for cavity fields. *Physical Review A*, 88(6):062324, 12 2013.

[61] D. L. McAuslan et al. Photon-echo quantum memories in inhomogeneously broadened two-level atoms. *Physical Review A*, 84(2):022309, 8 2011.

[62] M. Bonarota, J. Dajczgewand, A. Louchet-Chauvet, J.-L. Le Gouët, and T. Chanelière. Photon echo with a few photons in two-level atoms. *Laser Physics*, 24(9):094003, 9 2014.

[63] K. I. Gerasimov et al. Quantum memory in an orthogonal geometry of silenced echo retrieval. *Optics and Spectroscopy*, 123(2):211–216, 2017.

[64] J. Dajczgewand, J.-L. Le Gouët, A. Louchet-Chauvet, and T. Chanelière. Large efficiency at telecom wavelength for optical quantum memories. *Optics Letters*, 39(9):2711–4, 2014.

[65] K. R. Ferguson, S. E. Beavan, J. J. Longdell, and M. J. Sellars. Generation of Light with Multimode Time-Delayed Entanglement Using Storage in a Solid-State Spin-Wave Quantum Memory. *Physical Review Letters*, 117(2):1–5, 2016.

[66] S. Moiseev, and N. Arslanov. Efficiency and fidelity of photon-echo quantum memory in an atomic system with longitudinal inhomogeneous broadening. *Physical Review A*, 78(2):023803, 8 2008.

[67] B. C. Buchler, M. Hosseini, G. Hétet, B. M. Sparkes, and P. K. Lam. Precision spectral manipulation of optical pulses using a coherent photon echo memory. *Optics letters*, 35(7):1091–3, 4 2010.

[68] M. P. Hedges, J. J. Longdell, Y. Li, and M. J. Sellars. Efficient quantum memory for light. *Nature*, 465(7301):1052–6, 6 2010.

[69] G. Hétet, and D. Guéry-Odelin. Spin wave diffraction control and read-out with a quantum memory for light. *New Journal of Physics*, 17(7), 2015.

[70] G. Hétet et al. Photon echoes generated by reversing magnetic field gradients in a rubidium vapor. *Optics letters*, 33(20):2323–5, 10 2008.

[71] B. M. Sparkes, M. Hosseini, G. Hétet, P. K. Lam, and B. C. Buchler. ac Stark gradient echo memory in cold atoms. *Physical Review A*, 82(4):043847, 10 2010.

[72] S. a. Moiseev, and W. Tittel. Optical quantum memory with generalized time-reversible atomlight interaction. *New Journal of Physics*, 13(6):063035, 6 2011.

[73] M. R. Hush, A. R. R. Carvalho, M. P. Hedges, and M. R. James. Analysis of the operation of gradient echo memories using a quantum input-output model. *New Journal of Physics*, 15(8):085020, 8 2013.

[74] D. B. Higginbottom et al. Dual-rail optical gradient echo memory. *Optics express*, 23(19):24937–44, 2015.

[75] D. Higginbottom et al. Spatial-mode storage in a gradient-echo memory. *Physical Review A*, 86(2):1–10, 8 2012.

[76] Q. Glorieux, J. B. Clark, A. M. Marino, Z. Zhou, and P. D. Lett. Temporally multiplexed storage of images in a gradient echo memory. *Optics Express*, 20(11):12350, 5 2012.

[77] J. B. Clark, Q. Glorieux, and P. D. Lett. Spatially addressable readout and erasure of an image in a gradient echo memory. *New Journal of Physics*, 15, 2013.

[78] X. Zhang, A. Kalachev, and O. Kocharovskaya. Quantum storage based on control-field angular scanning. *Physical Review A*, 87(1):013811, 1 2013.

[79] A. Kalachev, and O. Kocharovskaya. Quantum storage via refractive-index control. *Physical Review A - Atomic, Molecular, and Optical Physics*, 83(5):1–6, 2011.

[80] J. Clark, K. Heshami, and C. Simon. Photonic quantum memory in two-level ensembles based on modulating the refractive index in time: Equivalence to gradient echo memory. *Physical Review A - Atomic, Molecular, and Optical Physics*, 86(1):1–5, 2012.

[81] A. E. Kozhekin, K. Molmer, K. M\olmer, and E. S. Polzik. Quantum memory for light. *Physical Review A*, 62(February):033809, 2000.

[82] M. Crisp. Propagation of Small-Area Pulses of Coherent Light through a Resonant Medium. *Physical Review A*, 1(6):1604–1611, 1970.

[83] K. F. Reim et al. Towards high-speed optical quantum memories. *Nature Photon.*, 4(4):218–221, 2010.

[84] K. F. Reim et al. Single-Photon-Level Quantum Memory at Room Temperature. *Physical Review Letters*, 107(5):053603, 7 2011.

[85] D. G. England et al. High-fidelity polarization storage in a gigahertz bandwidth quantum memory. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(12):124008, 2012.

[86] P. S. Michelberger et al. Interfacing GHz-bandwidth heralded single photons with a warm vapour Raman memory. *New Journal of Physics*, 17(4):043006, 4 2015.

[87] D. G. England et al. Storage and Retrieval of THz-Bandwidth Single Photons Using a Room-Temperature Diamond Quantum Memory. *Physical Review Letters*, 114(5):053602, 2 2015.

[88] M. R. Sprague et al. Broadband single-photon-level memory in a hollow-core photonic crystal fibre. *Nature Photonics*, 8(4):287–291, 2014.

[89] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414(6862):413–418, 11 2001.

[90] A. I. Lvovsky, B. C. Sanders, and W. Tittel. Optical quantum memory. *Nature Photonics*, 3(12):706–714, 12 2009.

[91] M. Parniak et al. Wavevector multiplexed atomic quantum memory via spatially-resolved single-photon detection. *Nature Communications*, 8(1):2140, 2017.

[92] Y. Jiang, J. Rui, X.-H. Bao, and J.-W. Pan. Dynamical zeroing of spin-wave momentum to suppress motional dephasing in an atomic-ensemble quantum memory. *Physical Review A*, 93(6):063819, 6 2016.

[93] A. Gorshkov, A. André, M. Lukin, and A. Sørensen. Photon storage in Λ-type optically dense atomic media. I. Cavity model. *Physical Review A*, 76(3):033804, 9 2007.

[94] D. J. Saunders et al. Cavity-Enhanced Room-Temperature Broadband Raman Memory. *Physical Review Letters*, 116(9):24–27, 2016.

[95] X.-H. Bao et al. Efficient and long-lived quantum memory with cold atoms inside a ring cavity. *Nature Physics*, 8(7):517–521, 2012.

[96] E. Bimbard et al. Homodyne Tomography of a Single Photon Retrieved on Demand from a Cavity-Enhanced Cold Atom Memory. *Physical Review Letters*, 112(3):033601, 1 2014.

[97] J. Borregaard et al. Scalable photonic network architecture based on motional averaging in room temperature gas. *Nature Communications*, 7(7):11356, 4 2016.

[98] M. Hosseini et al. Coherent optical pulse sequencer for quantum applications. *Nature*, 461(7261):241–5, 9 2009.

[99] B. Sparkes et al. Precision Spectral Manipulation: A Demonstration Using a Coherent Optical Memory. *Physical Review X*, 2(2):021011, 6 2012.

[100] G. Campbell, M. Hosseini, B. M. Sparkes, P. K. Lam, and B. C. Buchler. Time- and frequency-domain polariton interference. *New Journal of Physics*, 14(3):033022, 3 2012.

[101] K. F. Reim et al. Multipulse Addressing of a Raman Quantum Memory: Configurable Beam Splitting and Efficient Readout. *Physical Review Letters*, 108(26):263602, 6 2012.

[102] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, 1 2001.

[103] G. T. Campbell et al. Configurable Unitary Transformations and Linear Logic Gates Using Quantum Memories. *Physical Review Letters*, 113(6):063601, 8 2014.

[104] S. Kröll, and U. Elman. Photon-echo-based logical processing. *Optics Letters*, 18(21):1834, 1993.

[105] B. Albrecht, P. Farrera, X. Fernandez-Gonzalvo, M. Cristiani, and H. de Riedmatten. A waveguide frequency converter connecting rubidium-based quantum memories to the telecom C-band. *Nature communications*, 5:3376, 2014.

[106] S. Zaske, A. Lenhard, and C. Becher. Efficient frequency downconversion at the single photon level from the red spectral range to the telecommunications C-band. *Optics Express*, 19(13):12825, 6 2011.

[107] R. Ikuta et al. Wide-band quantum interface for visible-to-telecommunication wavelength conversion. *Nature Communications*, 2(1):1544–1545, 2011.

[108] R. Ikuta et al. High-fidelity conversion of photonic quantum information to telecommunication wavelength with superconducting single-photon detectors. *Physical Review A - Atomic, Molecular, and Optical Physics*, 87(1):3–6, 2013.

[109] B. M. Sparkes. *Storage and Manipulation of Optical Information Using Gradient Echo Memory in Warm Vapours and Cold Ensembles Australian National*. PhD thesis, 2013.

[110] Y.-W. Cho et al. Highly efficient optical quantum memory with long coherence time in cold atoms. *Optica*, 3(1):100, 1 2016.

[111] D. A. Steck. "Rubidium 87 D Line Data," available online at http://steck.us/alkalidata. 2010.

[112] T. C. Ralph, P. K. Lam, and R. E. S. Polkinghorne. Characterizing teleportation in optics. *Journal of Optics B: Quantum and Semiclassical Optics*, 1(4):483–489, 8 1999.

[113] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler. Unconditional room-temperature quantum memory. *Nature Physics*, 7(10):794–798, 6 2011.

[114] D. A. Steck. "Cesium D Line Data," available online at http://steck.us/alkalidata. 2010.

[115] N. V. Corzo et al. Large Bragg Reflection from One-Dimensional Chains of Trapped Atoms Near a Nanoscale Waveguide. *Physical Review Letters*, 117(13):133603, 9 2016.

[116] a. André, and M. Lukin. Manipulating Light Pulses via Dynamically Controlled Photonic Band gap. *Physical Review Letters*, 89(14):143602, 9 2002.

[117] M. Bajcsy, A. S. Zibrov, and M. D. Lukin. Stationary pulses of light in an atomic medium. *Nature*, 426(6967):638–41, 12 2003.

[118] A. W. Brown, and M. Xiao. Frequency detuning and power dependence of reflection from an electromagnetically induced absorption grating. *Journal of Modern Optics*, 52(16):2365–2371, 2005.

[119] a. André, M. Bajcsy, a. Zibrov, and M. Lukin. Nonlinear Optics with Stationary Pulses of Light. *Physical Review Letters*, 94(6):063902, 2 2005.

[120] J.-H. Wu, M. Artoni, and G. C. La Rocca. Stationary light pulses in cold thermal atomic clouds. *Physical Review A*, 82(1):013807, 7 2010.

[121] G. T. Campbell et al. Direct imaging of slow, stored and stationary EIT polaritons. *Quantum Science and Technology*, 2(3):034010, 2017.

[122] S. A. Moiseev, and B. S. Ham. Generation of entangled lights with temporally reversed photon wave packets. *Physical Review A*, 71(5):053802, 5 2005.

[123] S. Moiseev, and B. Ham. Quantum manipulation of two-color stationary light: Quantum wavelength conversion. *Physical Review A*, 73(3):033812, 3 2006.

[124] S. Moiseev, and B. Ham. Quantum control and manipulations of three-color stationary light. *Journal of the Korean Physical Society*, 48(4):540–545, 2006.

[125] C.-L. Cui, J. Wu, J.-W. Gao, Y. Zhang, and N. Ba. Double photonic bandgaps dynamically induced in a tripod system of cold atoms. *Optics express*, 18(5):4538–4546, 2010.

[126] X. Su, and B. Ham. Dynamic control of the photonic band gap using quantum coherence. *Physical Review A*, 71(1):1–5, 2005.

[127] F. Zimmer, A. André, M. Lukin, and M. Fleischhauer. Coherent control of stationary light pulses. *Optics Communications*, 264(2):441–453, 8 2006.

[128] S. E. Harris. Electromagnetically induced transparency with matched pulses. *Physical Review Letters*, 70(5):552–555, 1993.

[129] F. Zimmer, J. Otterbach, R. Unanyan, B. Shore, and M. Fleischhauer. Dark-state polaritons for multicomponent and stationary light fields. *Physical Review A*, 77(6):063823, 6 2008.

[130] K. Surmacz et al. Efficient spatially resolved multimode quantum memory. *Physical Review A - Atomic, Molecular, and Optical Physics*, 78(3):1–9, 2008.

[131] S. P. Tewari, and G. S. Agarwal. Control of phase matching and nonlinear generation in dense media by resonant fields. *Physical Review Letters*, 56(17):1811–1814, 1986.

[132] M. Jain, G. Y. Yin, J. E. Field, and S. E. Harris. Observation of electromagnetically induced phase matching. *Optics Letters*, 18(12):998, 1993.

[133] T. Peters, Y.-H. Chen, J.-S. Wang, Y.-W. Lin, and I. a. Yu. Observation of phase variation within stationary light pulses inside a cold atomic medium. *Optics Letters*, 35(2):151, 1 2010.

[134] Y.-f. Chen, Y.-m. Kao, W.-h. Lin, and I. A. Yu. Phase variation and shape distortion of light pulses in electromagnetically induced transparency media. *Physical Review A*, 74(6):063807, 12 2006.

[135] K. Hansen, and K. Mølmer. Trapping of light pulses in ensembles of stationary Λ atoms. *Physical Review A*, 75(5):053802, 5 2007.

[136] K. Hansen, and K. Mølmer. Stationary light pulses in ultracold atomic gases. *Physical Review A*, 75(6):065804, 6 2007.

[137] J.-H. Wu, M. Artoni, and G. C. La Rocca. Stationary light pulses in cold thermal atomic clouds. *Physical Review A*, 82(1):013807, 7 2010.

[138] T. Peters et al. Formation of stationary light in a medium of nonstationary atoms. *Physical Review A - Atomic, Molecular, and Optical Physics*, 85(2):023838, 2 2012.

[139] G. Nikoghosyan, and M. Fleischhauer. Stationary light in cold-atomic gases. *Physical Review A*, 80(1):013818, 7 2009.

[140] Y.-W. Lin et al. Stationary Light Pulses in Cold Atomic Media and without Bragg Gratings. *Physical Review Letters*, 102(21):213601, 5 2009.

[141] S. a. Moiseev, a. I. Sidorova, and B. S. Ham. Stationary and quasistationary light pulses in three-level cold atomic systems. *Physical Review A*, 89(4):043802, 4 2014.

[142] I. Iakoupov, J. R. Ott, D. E. Chang, and A. S. Sørensen. Dispersion relations for stationary light in one-dimensional atomic ensembles. *Physical Review A*, 94(5):053824, 11 2016.

[143] J. Otterbach, R. Unanyan, and M. Fleischhauer. Confining Stationary Light: Dirac Dynamics and Klein Tunneling. *Physical Review Letters*, 102(6):063602, 2 2009.

[144] J. Otterbach, J. Ruseckas, R. G. Unanyan, G. Juzeliūnas, and M. Fleischhauer. Effective Magnetic Fields for Stationary Light. *Physical Review Letters*, 104(3):033903, 1 2010.

[145] D. E. Chang et al. Crystallization of strongly interacting photons in a nonlinear optical fibre. *Nature Physics*, 4(11):884–889, 9 2008.

[146] M. Kiffner, and M. J. Hartmann. Master equation approach for interacting slow- and stationary-light polaritons. *Physical Review A*, 82(3):033813, 9 2010.

[147] M. Fleischhauer, J. Otterbach, and R. G. Unanyan. Bose-Einstein Condensation of Stationary-Light Polaritons. *Physical Review Letters*, 101(16):163601, 10 2008.

[148] P. Das, C. Noh, and D. G. Angelakis. Realization of the driven nonlinear Schrödinger equation with stationary light. *EPL (Europhysics Letters)*, 103(3):34001, 8 2013.

[149] S. a. Carvalho, and L. E. E. de Araujo. Electromagnetically induced blazed grating at low light levels. *Physical Review A*, 83(5):053825, 5 2011.

[150] M. Hafezi, D. E. Chang, V. Gritsev, E. Demler, and M. D. Lukin. Quantum transport of strongly interacting photons in a one-dimensional nonlinear waveguide. *Physical Review A - Atomic, Molecular, and Optical Physics*, 85:1–22, 2012.

[151] M. Hafezi, D. E. Chang, V. Gritsev, E. Demler, and M. D. Lukin. Quantum transport of strongly interacting photons in a one-dimensional nonlinear waveguide. *Physical Review A - Atomic, Molecular, and Optical Physics*, 85:1–20, 2012.

[152] Z. Ullah et al. Observation of the four wave mixing photonic band gap signal in electromagnetically induced grating. *Optics Express*, 22(24):29544, 11 2014.

[153] Y.-M. Liu, F. Gao, C.-H. Fan, and J.-H. Wu. Asymmetric light diffraction of an atomic grating with PT symmetry. *Optics Letters*, 42(21):4283, 11 2017.

[154] F. Blatt, L. S. Simeonov, T. Halfmann, and T. Peters. Stationary light pulses and narrowband light storage in a laser-cooled ensemble loaded into a hollow-core fiber. *Physical Review A*, 94(4):043833, 10 2016.

[155] G. R. Dennis, J. J. Hope, and M. T. Johnsson. XMDS2: Fast, scalable simulation of coupled stochastic partial differential equations. *Computer Physics Communications*, 184(1):201–208, 1 2013.

[156] A. Tranter et al. Multiparameter optimisation of a magneto-optical trap using deep learning (in review).

[157] B. He, and A. Scherer. Continuous-mode effects and photon-photon phase gate performance. *Physical Review A*, 85(3):033814, 3 2012.

[158] W. J. Munro, K. Nemoto, and T. P. Spiller. Weak nonlinearities: a new route to optical quantum computation. *New Journal of Physics*, 7:137–137, 5 2005.

[159] J. Barrat, and C. Cohen-Tannoudji. Élargissement et déplacement des raies de résonance magnétique causés par une excitation optique. *Journal de Physique et le Radium*, 22(7):443–450, 1961.

[160] K. M. Beck, M. Hosseini, Y. Duan, and V. Vuletić. Large conditional single-photon cross-phase modulation. *Proceedings of the National Academy of Sciences*, 113(35):9740–9744, 8 2016.

[161] O. Lahad, and O. Firstenberg. Induced Cavities for Photonic Quantum Gates. *Physical Review Letters*, 119(11):1–7, 2017.

[162] M. Sondermann, and G. Leuchs. The phase shift induced by a single atom in free space. *Journal of the European Optical Society: Rapid Publications*, 8:13052, 2013.

[163] V. Venkataraman, K. Saha, and A. L. Gaeta. Phase modulation at the few-photon level for weak-nonlinearity-based quantum computing. *Nature Photonics*, 7(2):138–141, 2013.

[164] N. Sinclair et al. Proposal and proof-of-principle demonstration of non-destructive detection of photonic qubits using a Tm:LiNbO3waveguide. *Nature Communications*, 7:1–6, 2016.

[165] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble. Measurement of Conditional Phase Shifts for Quantum Logic. *Physical Review Letters*, 75(25):4710–4713, 12 1995.

[166] M. J. Werner, and A. Imamoğlu. Photon-photon interactions in cavity electromagnetically induced transparency. *Physical Review A*, 61(1):011801, 12 1999.

[167] J. Volz, M. Scheucher, C. Junge, and A. Rauschenbeutel. Nonlinear π phase shift for single fibre-guided photons interacting with a single resonator-enhanced atom. *Nature Photonics*, 8(12):965–970, 12 2014.

[168] H. Schmidt, and A. Imamoglu. Giant Kerr nonlinearities obtained by electromagnetically induced transparency. *Optics Letters*, 21(23):1936, 12 1996.

[169] S. Harris, and L. Hau. Nonlinear Optics at Low Light Levels. *Physical Review Letters*, 82:4611–4614, 1999.

[170] S. E. Harris, and Y. Yamamoto. Photon switching by quantum interference. *Physical Review Letters*, 81(17):3611–3614, 1998.

[171] D. A. Braje, V. Balić, G. Y. Yin, and S. E. Harris. Low-light-level nonlinear optics with slow light. *Physical Review A*, 68(4):041801, 2003.

[172] V. Venkataraman, K. Saha, P. Londero, and A. L. Gaeta. Few-Photon All-Optical Modulation in a Photonic Band-Gap Fiber. *Physical Review Letters*, 107(19):193902, 11 2011.

[173] H. Kang, and Y. Zhu. Observation of large Kerr nonlinearity at low light intensities. *Physical review letters*, 91(August):093601, 2003.

[174] H.-Y. Lo et al. Electromagnetically-induced-transparency-based cross-phase-modulation at attojoule levels. *Physical Review A*, 83(4):041804, 4 2011.

[175] B.-W. Shiau, M.-C. Wu, C.-C. Lin, and Y.-C. Chen. Low-Light-Level Cross-Phase Modulation with Double Slow Light Pulses. *Physical Review Letters*, 106(19):193006, 5 2011.

[176] Z.-y. Liu et al. Large Cross-Phase Modulations at the Few-Photon Level. *Physical Review Letters*, 117(20):203601, 11 2016.

[177] A. Feizpour, M. Hallaji, G. Dmochowski, and A. M. Steinberg. Observation of the nonlinear phase shift due to single post-selected photons. *Nature Physics*, 11(11):905–909, 8 2015.

[178] M. V. Pack, R. M. Camacho, and J. C. Howell. Transients of the electromagnetically-induced-transparency-enhanced refractive Kerr nonlinearity. *Physical Review A - Atomic, Molecular, and Optical Physics*, 76(3):1–13, 2007.

[179] G. F. Sinclair. Time-dependent cross-phase-modulation in Rb87. *Physical Review A - Atomic, Molecular, and Optical Physics*, 79(2), 2009.

[180] L. Deng, M. G. Payne, and W. R. Garrett. Electromagnetically-induced-transparency-enhanced Kerr nonlinearity: Beyond steady-state treatment. *Physical Review A. Atomic, Molecular, and Optical Physics*, 64(2):1–8, 2001.

[181] C. Ottaviani, S. Rebić, D. Vitali, and P. Tombesi. Quantum phase-gate operation based on nonlinear optics: Full quantum analysis. *Physical Review A*, 73(1):1–4, 1 2006.

[182] S. Rebić, C. Ottaviani, G. Di Giuseppe, D. Vitali, and P. Tombesi. Assessment of a quantum phase-gate operation based on nonlinear optics. *Physical Review A - Atomic, Molecular, and Optical Physics*, 74(3):1–11, 2006.

[183] M. V. Pack, R. M. Camacho, and J. C. Howell. Transients of the electromagnetically-induced-transparency-enhanced refractive Kerr nonlinearity. *Physical Review A - Atomic, Molecular, and Optical Physics*, 76(3):1–9, 2007.

[184] A. Feizpour, G. Dmochowski, and A. M. Steinberg. Short-pulse cross-phase modulation in an electromagnetically-induced-transparency medium. *Physical Review A*, 93(1):013834, 6 2014.

[185] G. Dmochowski et al. Experimental Demonstration of the Effectiveness of Electromagnetically Induced Transparency for Enhancing Cross-Phase Modulation in the Short-Pulse Regime. *Physical Review Letters*, 116(17):173002, 4 2016.

[186] M. Hosseini, B. M. Sparkes, G. T. Campbell, P. K. Lam, and B. C. Buchler. Storage and manipulation of light using a Raman gradient-echo process. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(12):124004, 6 2012.

[187] I. Friedler, G. Kurizki, and D. Petrosyan. Deterministic quantum logic with photons via optically induced photonic band gaps. *Physical Review A - Atomic, Molecular, and Optical Physics*, 71(2):1–8, 2005.

[188] Y.-H. Chen et al. Demonstration of the Interaction between Two Stopped Light Pulses. *Physical Review Letters*, 108(17):173603, 4 2012.

[189] M. D. Lukin, and a. Imamoğlu. Controlling photons using electromagnetically induced transparency. *Nature*, 413(6853):273–6, 9 2001.

[190] A. K. Mohapatra, T. R. Jackson, and C. S. Adams. Coherent optical detection of highly excited rydberg states using electromagnetically induced transparency. *Physical Review Letters*, 98(11):1–4, 2007.

[191] J. D. Pritchard et al. Cooperative Atom-Light Interaction in a Blockaded Rydberg Ensemble. *Physical Review Letters*, 105(19):193603, 11 2010.

[192] V. Parigi et al. CLEO®/Europe - IQEC 2013 Observation and measurement of interaction-induced dispersive optical nonlinearities in an ensemble of cold Rydberg atoms. *2013 Conference on Lasers and Electro-Optics Europe and International Quantum Electronics Conference, CLEO/Europe-IQEC 2013*, 233602(DECEMBER):1–5, 2013.

[193] Y. M. Hao et al. Quantum controlled-phase-flip gate between a flying optical photon and a Rydberg atomic ensemble. *Scientific Reports*, 5(November 2014):1–7, 2015.

[194] S. Das et al. Photonic controlled- phase gates through Rydberg blockade in optical cavities. *Physical Review A*, 93(4):1–6, 2016.

[195] I. Iakoupov, J. Borregaard, and A. S. Sørensen. Controlled-phase Gate for Photons Based on Stationary Light. *Physical Review Letters*, 120(1):010502, 1 2018.

[196] J. Shapiro. Single-photon Kerr nonlinearities do not help quantum computation. *Physical Review A*, 73(6):062305, 6 2006.

[197] J. H. Shapiro, and M. Razavi. Continuous-time cross-phase modulation and quantum computation. *New Journal of Physics*, 9, 2007.

[198] J. Gea-Banacloche. Impossibility of large phase shifts via the giant Kerr effect with single-photon wave packets. *Physical Review A*, 81(4):043823, 4 2010.

[199] D. J. Brod, J. Combes, and J. Gea-Banacloche. Two photons co- and counterpropagating through N cross-Kerr sites. *Physical Review A*, 94(2):023833, 8 2016.

[200] D. J. Brod, and J. Combes. Passive CPHASE Gate via Cross-Kerr Nonlinearities. *Physical Review Letters*, 117(8):080502, 8 2016.

[201] B. He, A. MacRae, Y. Han, A. I. Lvovsky, and C. Simon. Transverse multimode effects on the performance of photon-photon gates. *Physical Review A*, 83(2):022312, 2 2011.

[202] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, 6 1982.

[203] T. Toffoli. Reversible computing. *Automata, Languages and Programming. ICALP 1980. Lecture Notes in Computer Science*, 85:632–644, 1980.

[204] E. Fredkin, and T. Toffoli. Conservative Logic. *International Journal of Theoretical Physics*, 21(3/4):219–253, 1982.

[205] T. Sleator, and H. Weinfurter. Realizable Universal Quantum Logic Gates. *Physical Review Letters*, 74(20):4087–4090, 5 1995.

[206] P. W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509, 10 1997.

[207] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96*, pages 212–219, New York, New York, USA, 1996. ACM Press.

[208] H. J. Kimble. The quantum internet. *Nature*, 453(7198):1023–1030, 2008.

[209] I. L. Chuang, and Y. Yamamoto. Simple quantum computer. *Physical Review A*, 52(5):3489–3496, 11 1995.

[210] B. Hacker, S. Welte, G. Rempe, and S. Ritter. A photon-photon quantum gate based on a single atom in an optical resonator. *Nature*, pages 1–6, 2016.

[211] L. M. Duan, and H. J. Kimble. Scalable photonic quantum computation through cavity-assisted interactions. *Physical Review Letters*, 92(12):1–4, 2004.

[212] G. J. Milburn. Quantum optical Fredkin gate. *Physical Review Letters*, 62(18):2124–2127, 5 1989.