

**Towards storage and retrieval of non-classical light
in a broadband quantum memory
- An investigation of free-space and cavity Raman
memories**

Theresa F. M. Champion
Christ Church, Oxford



Submitted for the degree of Doctor of Philosophy
Trinity Term 2015

Supervised by
Prof. Ian A. Walmsley

Clarendon Laboratory
University of Oxford
United Kingdom

**Towards storage and retrieval of non-classical light
in a broadband quantum memory
- An investigation of free-space and cavity Raman memories**

Theresa F. M. Champion
Christ Church, Oxford

Submitted for the degree of Doctor of Philosophy
Trinity Term 2015

Abstract

Photonic quantum information processing has emerged as a powerful platform for realising quantum-enhanced technologies. In order to be scalable, many of these technologies depend on the availability of a suitable quantum memory for the coherent storage and on-demand retrieval of photonic quantum states.

In this thesis, I investigate broadband light storage in a room-temperature Raman memory, implemented both in free space and, for the first time, inside a low-finesse optical cavity designed for low-noise operation.

The ability of the Raman memory to preserve phase coherence was tested by storing coherent polarisation states in two spatially separate atomic ensembles. Polarisation storage with a fidelity of up to $97 \pm 1\%$ was demonstrated by performing full process tomography on the system.

The Raman memory was then interfaced for the first time with a spontaneous parametric downconversion (SPDC) source of heralded, GHz-bandwidth single photons. The memory performance was characterised by measuring the second-order autocorrelation of the retrieved fields. While the SPDC input photon statistics showed a clear influence on the statistics of the retrieved field, four-wave mixing (FWM) noise, stimulated by spontaneous Raman scattering, prevented the preservation of non-classical photon statistics during read-out.

Suppressing this source of noise represents the last remaining challenge for realising a broadband single-photon Raman memory suitable for quantum information applications. To this end, I demonstrate a novel cavity implementation of the Raman memory which reduces the FWM contribution relative to the signal field by re-distributing the density of states into which the noise photons can be scattered. Cavity-enhanced memory operation was investigated using weak coherent input states, showing a significant improvement of the signal-to-noise ratio compared to the free-space memory implementation. This proof-of-principle demonstration suggests that cavity Raman memories may offer a practical route towards low-noise, high-bandwidth quantum storage at room temperature.

To my parents

Acknowledgements

I have been tremendously lucky to have worked alongside so many wonderful people over the last five years. Their friendship, sense of humour, camaraderie and unwavering support have carried me across some of the most exciting and challenging times. To have been part of such an exceptional team is something I will always treasure.

First of all, I owe an enormous debt of gratitude to Ian Walmsley for his brilliant guidance, foresight and good-humoured equanimity. I feel incredibly fortunate to have been allowed to pursue my research work under his excellent supervision.

I could not imagine the Ultrafast group without Josh Nunn; from the very first day, which – by good fortune – was mostly spent lunching at Chiang Mai, his inimitable humour, warmth and utterly self-deprecating intellect suffused everything with humanity.

It was an enormous pleasure to work with Duncan England who taught me the ropes and imbued the lab with a wonderful sense of conviviality. His love of cricket may not have quite rubbed off on me, but I think I can still faintly hear the unfathomable endlessness of the test series broadcast. Throughout the years, Patrick Michelberger's ceaseless energy, rigour and sheer tenacity – especially when faced with the endurance test that were autocorrelation measurements – were second to none.

When embarking on the adventure of the cavity memory, I could not have hoped for a better partner in crime than Dylan Saunders whose amazing spirit, resourcefulness and ingenuity were inspirational. I also owe him a huge thank you for helping me keep calm during these last few stressful months of writing. I am immensely grateful to Joseph Munns for all of his manifold contributions; his outstanding work and clear-mindedness will surely lead the cavity memory to new heights. It was also such good fortune to have Patrick Ledingham and Cheng Qiu join the team – their great skill, insight and dedication made all the difference. I am so happy to be able to leave the cavity memory in the very able hands of such a marvellous team and will be eagerly awaiting news of their next achievements.

Over the course of many years, the Raman memory, in its various incarnations, has brought together a terrific group of fellow students and post-docs: Klaus Reim's pioneering work laid the foundation for what was to follow, KC Lee and Michael Sprague were both awesome examples of understated brilliance, Xian-Min Jin and Steve Kolthammer made indispensable contributions right down to the most hands-on help with the experiment, and Kris Kaczmarek's remarkable proficiency is only equalled by his infectious enthusiasm.

I would also like to thank Eilon Poem-Kalogerakis, Benjamin Brecht and Amir Feizpour for their extremely helpful comments at various stages of drafting this thesis.

To my friends at college and beyond who have enriched every aspect of my time at Oxford, to Ray whose unconditional love and support mean the world to me, and finally, to my parents and grandmother – thank you for more than I can possibly say.

Author's publications

Journal publications

1. P. S. Michelberger, T. F. M. Champion, M. R. Sprague, K. Kaczmarek, M. Barbieri, X.-M. Jin, D. G. England, W. S. Kolthammer, D. Saunders, J. Nunn and I. A. Walmsley. “**Interfacing GHz-Bandwidth Single Photons with a Room-temperature Raman Quantum Memory**” *New Journal of Physics*, **17**, 043006 (2015).
2. J. Nunn, N. K. Langford, W. S. Kolthammer, T. F. M. Champion, M. R. Sprague, P. S. Michelberger, X.-M. Jin, D. G. England, and I. A. Walmsley. “**Enhancing Multiphoton Rates with Quantum Memories**” *Phys. Rev. Lett.*, **110**, 133601 (2013).
3. D. G. England, P. S. Michelberger, T. F. M. Champion, K. F. Reim, K. C. Lee, M. R. Sprague, X.-M. Jin, N. K. Langford, W. S. Kolthammer, J. Nunn and I. A. Walmsley. “**High-Fidelity Polarization Storage in a Gigahertz Bandwidth Quantum Memory**” *Journal of Physics B*, **45**, 124008 (2012).
4. K. F. Reim, J. Nunn, X.-M. Jin, P. S. Michelberger, T. F. M. Champion, D. G. England, K. C. Lee, N. K. Langford and I. A. Walmsley. “**Multipulse Addressing of a Raman Quantum Memory: Configurable Beam Splitting and Efficient Readout**” *Phys. Rev. Lett.*, **108**, 263602 (2012).

Conferences

1. T. F. M. Champion, J. Munns, C. Qiu, P. M. Ledingham, J. Nunn, D. Saunders and I. A. Walmsley. “**Bad Cavities for Good Memories: Storing Broadband Photons with Low Noise**” *Conference on Lasers and Electro-Optics (CLEO)*, May 2015.
2. P. S. Michelberger, T. F. M. Champion, M. R. Sprague, K. Kaczmarek, M. Barbieri, X.-M. Jin, D. G. England, W. S. Kolthammer, D. Saunders, J. Nunn and I. A. Walmsley. “**Interfacing GHz-Bandwidth Single Photons with a Room-temperature Raman Quantum Memory**” *Conference on Quantum Communication, Measurement & Computing (QCMC)*, November 2014.
3. T. F. M. Champion, P. S. Michelberger, D. G. England, K. F. Reim, M. R. Sprague, K. C. Lee, X.-M. Jin, N. K. Langford, W. S. Kolthammer, J. Nunn and I. A. Walmsley. “**Broadband Quantum Memory in Caesium**” *Quantum Coherence Workshop*, March 2013.
4. D. G. England, P. S. Michelberger, T. F. M. Champion, K. F. Reim, K. C. Lee, M. R. Sprague, X.-M. Jin, M. Barbieri, W. S. Kolthammer, N. K. Langford, J. Nunn and I. A. Walmsley. “**High-Fidelity Polarization Storage in a Broadband Quantum Memory**” *Spring Meeting of the German Physical Society*, March 2012.

Contents

Acknowledgements	iii
Author's publications	v
List of Figures	xi
1 Introduction	1
1.1 Quantum memory	3
1.2 Applications	4
1.2.1 Multiplexing strategies	5
1.2.2 Quantum communication	8
1.2.3 Quantum computing	10
1.3 Figures of merit	11
1.3.1 Fidelity	12
1.3.2 Memory efficiency	13
1.3.3 Storage time	14
1.3.4 Bandwidth	15
1.3.5 Time-bandwidth product	15
1.3.6 Multi-mode capacity	16
1.4 Review of quantum memories	16
1.4.1 Optical delay lines	17
1.4.2 Adiabatic memories	18
1.4.3 Photon-echo memory	24
1.5 Thesis outline	29
2 Theoretical overview of the Raman memory	32
2.1 Equations of motion	33
2.1.1 The light-matter Hamiltonian	33
2.1.2 Optical fields	35
2.1.3 The dipole interaction	36
2.1.4 Optical Bloch equations	38
2.1.5 Signal field propagation	40
2.1.6 Raman memory equations of motion	41

2.2	The Raman memory interaction	43
2.2.1	Storage and retrieval	43
2.2.2	Adiabatic limit	46
2.2.3	Raman memory coupling	48
2.2.4	Beam splitter interaction	49
2.3	Summary of Raman memory theory	52
3	The Raman memory experiment	53
3.1	Experimental setup	54
3.1.1	Storage medium	54
3.1.2	Laser system	57
3.1.3	Frequency lock	57
3.1.4	Signal and control field preparation	58
3.1.5	Optical pumping	62
3.1.6	Filtering and detection	63
3.2	Previous results	65
3.2.1	Storage of bright coherent states	65
3.2.2	Storage at the single-photon level	67
3.2.3	Multi-pulse addressing	71
4	Polarisation storage	74
4.1	Storage of polarisation states	75
4.1.1	Polarisation selection rules	75
4.2	The polarisation interferometer	76
4.3	Quantum process tomography	78
4.4	Experimental procedure	80
4.5	Results of the process tomography	83
5	Single photon storage	86
5.1	Single photon source	87
5.1.1	SHG and SPDC	87
5.1.2	Herald filtering and detection	90
5.2	Interfacing the memory with SPDC photons	91
5.2.1	Re-building the Raman memory	91
5.2.2	Feed-forward operation of the memory	95
5.2.3	Data acquisition	96
5.2.4	Temporal and spatial overlap	97
5.3	Memory operation	98
5.3.1	Heralding efficiency	98
5.3.2	Memory efficiency	99
5.3.3	Signal-to-noise ratio	102
5.3.4	$g^{(2)}$ -autocorrelation measurements	103
5.3.5	Results of the $g^{(2)}$ -measurements	106
5.4	Modelling FWM noise	108
5.4.1	Coherent interaction model	109
5.4.2	Incoherent interaction model	113

5.5	Further noise investigation	116
5.5.1	Build-up of FWM noise	116
5.5.2	Temperature dependence	117
5.5.3	FWM coupling strength	118
5.5.4	Spatial distribution of Stokes scattering	120
5.6	Strategies for reducing FWM noise	121
5.6.1	Phase-matching condition	121
5.6.2	Selection rules	124
5.6.3	Density of states	126
6	The cavity Raman memory	127
6.1	Input-output theory	127
6.1.1	Field propagation equations in the space-time domain	128
6.1.2	Linearised response in the space-frequency domain	130
6.2	Cavity Raman memory equations	132
6.3	Cavity enhancement	137
6.4	Mode selectivity	139
6.5	Noise suppression	142
6.6	The ring cavity design	143
6.7	The lossy ring cavity	144
6.8	Extinction	147
6.9	Ring cavity geometry	150
6.9.1	Gaussian beam propagation	151
6.10	Choice of cavity parameters	157
6.10.1	Free-spectral range and bandwidth	158
6.10.2	Control field coupling	159
6.10.3	Cavity mirrors	162
6.11	Optical pumping inside the cavity	163
6.12	Hänsch-Couillaud locking of the cavity	164
7	Experimental implementation of the cavity Raman memory	168
7.1	The ring cavity	169
7.1.1	Optical components	169
7.1.2	Coatings	170
7.2	Birefringent crystal	172
7.2.1	Alignment of crystal axes	173
7.2.2	Temperature control	175
7.3	Vapour cell	176
7.3.1	Cell heating	176
7.3.2	Optical depth	178
7.4	Experimental setup	181
7.4.1	Temporal overlap of signal and control pulses	182
7.4.2	Focusing	183
7.4.3	Piezo-electric control of the cavity length	183
7.4.4	Alignment	185
7.5	FSR measurements	185

7.6	Cavity locking	186
7.6.1	Locking signal	186
7.6.2	Locking procedure	188
7.6.3	Cavity stability	189
7.7	Control field coupling	191
7.7.1	Non-degenerate cavity response	191
7.7.2	Degenerate cavity response	192
7.8	Memory operation	195
7.8.1	Control field suppression	196
7.8.2	Signal detection and data acquisition	197
7.8.3	Storage of weak coherent states	197
7.9	Noise suppression	200
8	Conclusion	208
8.1	Summary	209
8.1.1	Free-space Raman memory	209
8.1.2	Cavity Raman memory	211
8.2	Outlook	212
8.2.1	Single-photon storage with the cavity Raman memory	212
8.2.2	Multiphoton rate enhancement	213
8.2.3	Light-matter Hong-Ou-Mandel interference	213
	Bibliography	216
A	Second-order autocorrelation function	234
A.1	Fock states	235
A.2	Coherent states	236
A.3	Thermal states	237
B	The cavity stability condition	240
B.1	Ray transfer matrix	240
B.2	Stability analysis	241

List of Figures

1.1	Multi-photon rate enhancement via memory-based synchronisation	7
1.2	Illustration of DLCZ-type entanglement swapping	9
1.3	Electromagnetically induced transparency	19
1.4	The Raman memory protocol	22
1.5	Level diagram for controlled reversible inhomogeneous broadening	25
1.6	Illustration of the gradient echo memory protocol	27
1.7	Level diagram of the atomic frequency comb memory	28
2.1	Illustration of the beam-splitter-type interaction of the Raman memory	51
3.1	Hyperfine structure of the caesium D_2 line	55
3.2	Experimental set-up of the Raman memory	60
3.3	Illustration of pulse-picking for the preparation of signal and control fields	61
3.4	First experimental results for coherent-state storage	66
3.5	Interference visibility of stored and retrieved signals	67
3.6	Storage of coherent states at the single-photon level	68
3.7	Memory lifetime measurements	69
3.8	Multi-pulse retrieval	72
3.9	Shaped retrieval illustrating the configurability of a memory-based beam splitter network	73
4.1	Diagram of the polarisation interferometer	77
4.2	Process matrices for polarisation storage reconstructed via process tomography	83
4.3	Process fidelity and memory efficiency as a function of storage time	84
5.1	PPKTP waveguide used for spontaneous parametric downconversion (SPDC)	88
5.2	Set-up of the single photon source	89
5.3	Filtering and detection set-up for the SPDC herald photon	91
5.4	Temperature-dependence of noise counts for different levels of frequency filtering	94
5.5	Timing diagram illustrating the feed-forward operation of the memory	95
5.6	Interfacing SPDC photons and coherent states with the memory	98
5.7	Storage and retrieval of heralded single photons	100
5.8	Hanbury Brown-Twiss set-up for $g^{(2)}$ -autocorrelation measurements	103

5.9	Results of the $g^{(2)}$ -autocorrelation experiments	106
5.10	Schematic of spontaneous anti-Stokes scattering and four-wave mixing (FWM)	109
5.11	Predictions of the incoherent interaction model	115
5.12	Build-up of FWM noise	116
5.13	Temperature dependence of memory efficiency and noise	117
5.14	Memory efficiency and noise as a function of control pulse energy	118
5.15	Test set-up for angle-tuning the signal and control beams	121
5.16	FWM phase matching condition	122
5.17	Proposed scheme for suppressing FWM noise based on selection rules	124
5.18	Raman memory coupling via two intermediate states	125
6.1	Diagram of ring cavity illustrating phase convention used in input-output theory	128
6.2	Density of states inside the cavity	131
6.3	Predicted $g^{(2)}$ -autocorrelation for varying ratios of memory and FWM coupling strengths	142
6.4	Diagram of two-junction ring cavity	145
6.5	Intensity spectra of a cavity with non-zero loss	147
6.6	Cavity extinction ratio as a function of output mirror reflectivity	150
6.7	Geometry of triangular ring cavity	150
6.8	Ray transfer matrix analysis of the unfolded ring cavity	153
6.9	Variation of the beam radius and radius of curvature for one round-trip of the cavity	154
6.10	Beam radius and radius of curvature as a function of mirror curvature	155
6.11	Stability diagram of the ring cavity	156
6.12	The effect of imperfect focusing on the beam radius variation	156
6.13	The effect of imperfect focusing on the radius of curvature	157
6.14	Normalised spectrum of a cavity satisfying the free-spectral-range condition	159
6.15	Normalised transmission spectrum of a doubly resonant cavity for orthogonal linear polarisations	160
6.16	Temperature dependence of path length change induced by intracavity birefringence	162
6.17	Peak transmission and resonance bandwidth as a function of output mirror reflectivity	163
6.18	Illustration of the Hänsch-Couillaud locking scheme	165
6.19	Predicted locking signal	167
7.1	Reflectance of dielectric mirror coatings as a function of wavelength	171
7.2	Test set-up used to determine alignment of birefringent crystal axes	173
7.3	Alignment of the axes of the birefringent BBO-crystal	174
7.4	Temperature control of the birefringence	175
7.5	Set-up used for measurements of the optical depth	180
7.6	Estimated optical depth for different temperatures	181
7.7	Experimental set-up of the cavity memory	182
7.8	Calibration of the piezo-electric ring actuator	184
7.9	Measured transmission spectrum of the cavity	186

7.10	Cavity stability test	190
7.11	Non-degenerate cavity response for orthogonal polarisations	192
7.12	Frequency scan of the degenerate cavity response	195
7.13	Storage and retrieval of weak coherent states.	197
7.14	Cavity Raman memory performance in terms of lifetime and dependence on control field power	199
7.15	Comparison of noise build-up in free-space and cavity implementations	201
7.16	Dependence of noise on control pulse energy	205

Chapter 1

Introduction

“The point is no longer that quantum mechanics is an extraordinarily (and for Einstein, unacceptably) peculiar theory, but that the world is an extraordinarily peculiar place.”

N. David Mermin

The properties of light have held an inexorable fascination for scientists since the earliest beginnings of empirical thought and inquiry. After more than two millennia of steadily refined theories of classical optics¹, the dawn of the twentieth century brought about a paradigmatic shift in our understanding of light and its interaction with matter.

The principles of quantum mechanics radically transformed the naturally intuitive perception of a world rooted in determinism and local realism^[1]. Replaced with a fundamentally probabilistic description of nature, these tenets gave way to profoundly new concepts, including superposition and entanglement. Quantum-enhanced technologies aim to harness quantum-mechanical properties in a way which promises to revolutionise the field of in-

¹Arguably the most significant contribution to the development of classical optics since Greek antiquity was made by the 11th-century scholar and polymath Ibn al-Haytham with his seven-volume treatise *Book of Optics*. Ibn al-Haytham’s evidence-based approach and willingness to question traditional doctrine furthermore anticipated the modern scientific method.

formation processing. Aptly described by Nicolas Gisin as “the science of turning quantum conundrums into potentially useful applications”^[2], *quantum information processing* (QIP) encompasses a wealth of potentially groundbreaking technologies. These include quantum computing, quantum metrology and quantum key distribution – the progenitor of large-scale quantum communication networks.

The vast potential of QIP applications has stimulated tremendous research efforts into the development of different physical systems and architectures^[3]. Amongst these, photonics has emerged as an alluringly powerful platform^[4]. The quantum state of light is an especially compelling system for encoding quantum information due to the unparalleled suitability for long-range distribution. Remarkable progress has been made across all component parts of a future photonic quantum network – from the development of sources for entangled qubits^[5–7], over integrated photonic circuits^[8–10] to sophisticated single-photon detectors^[11–13].

The practical realisation of many photonic QIP protocols depends on the ability to *store* quantum information encoded in light. This ability was recognised as key to overcoming a number of critical challenges: generating multiple simultaneous photons from probabilistic single-photon sources; extending quantum communication channels over long distances beyond the limits set by photon loss; synchronising independent, probabilistic processes in order to render composite protocols scalable. Over the last decade, these tasks have motivated the development of a suitable device – a *quantum memory* – which is capable of coherent storage and on-demand retrieval of optical quantum states. We will review the above applications more closely in the following sections.

In this thesis, I will examine a particular type of quantum memory based on the far-off-resonant Raman interaction in a room-temperature atomic ensemble. In contrast to

all-optical quantum memories (see Section 1.4.1), light storage in an atomic medium is based on the dynamic modulation of the interaction between light and matter. This highlights the conflicting requirements imposed on photonic qubits; while weakly interacting photons are ideal for the transmission of quantum information over long distances, atomic quantum memories need to induce strong, yet reversible coupling between the photons and the memory medium.

In the next section we will begin by distinguishing quantum memories from their classical counterpart. This is followed by an overview of some of the most important applications of quantum memories to quantum information processing (Section 1.2). In Section 1.3 we will consider the figures of merit commonly used to assess quantum memories and in Section 1.4 we will review different quantum memory schemes, as well as some recent experimental achievements. The final part of this chapter concludes with an outline of the structure and content of my thesis.

1.1 Quantum memory

Fundamentally, an atomic quantum memory acts as a light-matter interface for dynamically coupled optical and material modes. Photonic quantum information is reversibly mapped onto a stationary atomic excitation and faithfully stored for a controllable amount of time, before being re-converted into an optical output mode.

Unlike classical memories, which combine measurement and re-production, a quantum memory must not localise a quantum state through observation. It thereby evades irreversible wave-function collapse and loss of coherence which would be the inevitable result of measuring the quantum state. This fundamental difference is a result of the no-cloning

theorem which was developed in 1982 by Wootters and Zurek^[14] and Dieks^[15]. The theorem states that arbitrary unknown quantum states cannot be replicated or amplified owing to the linearity of quantum mechanics.

The no-cloning theorem has far-reaching consequences for QIP applications beyond the challenge of quantum state storage. For example, it precludes the use of classical error correction which takes advantage of informational redundancy by producing multiple copies of the encoded state. Instead, quantum error-correcting algorithms are needed to circumvent the no-cloning theorem and allow fault-tolerant quantum computation^[16–19]. In Section 1.2.2 we will consider a further implication of this theorem pertaining to long-distance quantum communication. The distribution of entangled quantum states suffers from exponential loss which severely limits the scalability of communication channels^[20]. This problem can be overcome by introducing a system of *quantum repeaters* for which quantum memories are an integral component^[21].

1.2 Applications

The study of quantum memories is of intrinsic interest; not only does it offer insight into the coherent interaction between light and matter, but it also elucidates the ways in which this interaction can be dynamically modulated and controlled. Beyond this foundational impetus, the development of quantum memories has been galvanised by the host of quantum information applications which – through the use of such memories – are brought within reach of realisation. The examples below provide an overview of some of these applications.

1.2.1 Multiplexing strategies

The probabilistic nature of many photonic quantum information primitives and protocols poses a severe obstacle for the scalability of these technologies. The compound probability of non-deterministic components rapidly vanishes, as we approach large-scale applications. Therefore the amount of physical resources required scales exponentially with the size of the system^[22], unless a multiplexing strategy is used. For example, this strategy could take the form of spatial or temporal multiplexing, or indeed some combination of both.

A spatially multiplexed system involves active switching with an array of resources operated in parallel. For instance, a deterministic source of single photons can be approximated by optically switching between multiple probabilistic sources^[23–26]. The emergence of integrated photonic circuits, such as monolithic sources of entangled photons^{[6] [27–30]}, offers promising platforms for realising spatially multiplexed systems.

Temporal multiplexing, in contrast, is based on the active synchronisation of probabilistic primitives which are repeatedly operated in series^[31–33]. Synchronisation is achieved by storing the heralded outputs of successful operations inside a quantum memory. On-demand recall, conditional on the success of all remaining processes of the composite photonic system, enables a repeat-until-success strategy which dramatically improves scalability. As an example of this strategy, we will consider time-multiplexed single-photon sources in the following section.

In addition, the method of spectral multiplexing applicable to frequency-encoded quantum information has been proposed by Sinclair *et al.*^[34]. Photons with distinct spectra, but equal temporal profiles, are simultaneously stored in a spectrally multimode memory (see Section 1.4.3). Selective retrieval of a particular frequency qubit is then achieved through a combination of adjustable frequency shifts and spectral filtering, based on the heralding

outcome of a preceding processing step. Sinclair *et al.* furthermore suggested that this approach involving frequency-multiplexed states could be combined with spatial and temporal degrees of freedom to achieve a multiplicative increase in the number of multiplexed modes.

Multi-photon rate enhancement

The generation of large numbers of single photons as a resource for sophisticated quantum information protocols has become an increasingly important challenge in the wake of rapid advances in the study of quantum-enhanced technologies^[9] [35–37]. Beyond the limited requirements of proof-of-principle demonstrations, these technologies necessitate readily attainable multi-photon states in order to achieve scalability.

The most common source of heralded single photons is based on the probabilistic process of spontaneous parametric downconversion (SPDC)^[38–40]. While suitable for small-scale experiments, these sources quickly approach the limits of feasibility if multiple simultaneous photons are required. In order to reduce the likelihood of unwanted higher-order photon pairs, SPDC sources have to be operated in a regime of low pair generation rates. This dramatically reduces the probability for simultaneous generation of photons from multiple sources operated in parallel. For instance, the current record of eight-photon states from four SPDC sources required measurement times of 40 hours to accumulate statistics^[41]. More generally, the probability of producing N single photons simultaneously scales exponentially as p^N , where $p \ll 1$ is the probability of generating and heralding a single photon². This scalability problem can be mitigated by interfacing each of the N sources with a quantum memory, as illustrated in Figure 1.1 (a). Once $N - 1$ heralded single photons have been stored, the memories are simultaneously read out, conditional on single-photon emission

²The probability p is typically $\sim 1\%$.

from the N^{th} source. The resulting synchronisation of the probabilistic sources enhances the N -fold coincidence rate by a factor which depends on the time-bandwidth product³ B as well as the memory efficiency η .

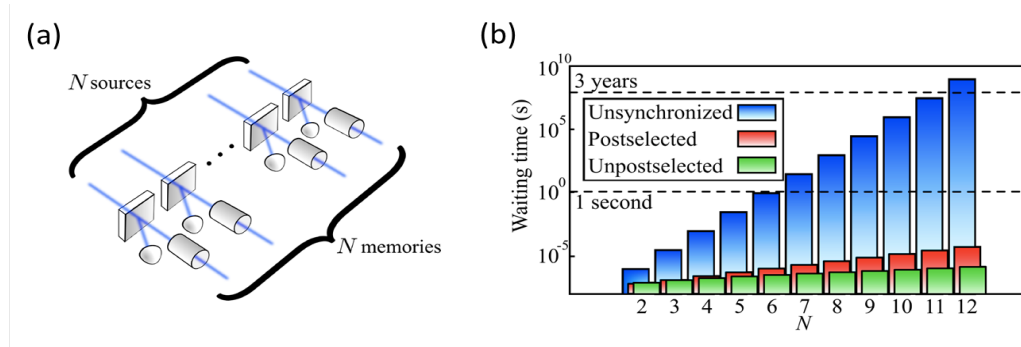


Figure 1.1: From reference^[42]: (a) Synchronisation of N heralded single-photon sources, each interfaced with a quantum memory. (b) Average waiting time for N -photon events with a threshold fidelity of 90% using unsynchronised (blue bars) and multiplexed downconversion sources with highly efficient memories (green bars), where storage and retrieval efficiencies of 99% have been assumed. If the number of photons is post-selected, the efficiency requirement can be relaxed to 75% (red bars).

As shown in reference^[42], the probability of generating N single photons in coincidence is given by

$$P_N = p^N \left(1 + \frac{(1-R)(1-p)\eta B}{1 + (B-1)[R(1-2p) + p]} \right)^N, \quad (1.1)$$

where R is the probability that the system is ready to be read out, i.e. the probability that $N-1$ memories have been charged and a photon is emitted from the N^{th} source. In the limit $\{RB, pB\} \ll 1$, the coincidence probability can be approximated as $P_N \approx (p\eta B)^N$. This scaling reflects the probability that a photon is produced in one of B discrete time bins afforded by the duration of the memory lifetime. For a given memory efficiency η , B therefore determines the number of trials with which a heralded photon can be synchronised with a previously charged memory before the stored excitation is lost to decoherence. In the opposite regime of $\{RB, pB\} \gg 1$, the coincidence rate eventually saturates to $P_N \approx \eta^N$,

³ B is defined as the product of the memory bandwidth and the memory lifetime (see Section 1.3.4).

independent of B .

An experimental demonstration of the synchronisation of two independent single-photon sources was reported in^[31]. It was furthermore shown that these synchronised sources could greatly enhance the generation of entangled photon pairs, compared to probabilistic sources.

The potential enhancement of multi-photon rates achievable with a system of synchronised sources is shown in Figure 1.1 (b). For instance, the average waiting time necessary for 12-photon events is reduced from more than three years to $\sim 100 \mu\text{s}$, assuming memories with $B = 1000$ and $\eta = 75\%$. Here, post-selection on the number of output photons compensates for storage and retrieval inefficiencies. The most critical memory parameter for this application is therefore the time-bandwidth product, rather than the absolute memory lifetime. This example demonstrates the ability of memory-based synchronisation protocols to ameliorate the scalability problem.

1.2.2 Quantum communication

The promise of unconditional security⁴ has inspired remarkable developments in the field of quantum communication. Seminal work by Bennett and Brassard^[44] and Ekert^[45] in 1984 and 1991, respectively, originated the study of quantum key distribution (QKD). While both QKD proposals are based on the transmission of photonic quantum states between remote parties, the Ekert protocol takes advantage of the non-classical correlations of entangled states. A random encryption key shared between Alice and Bob cannot be intercepted without perturbing the transmitted state^[46], thus enabling secure communication.

For long-distance quantum communication, photons are the only viable information carriers due to their unparalleled propagation speeds and weak interactions with the environ-

⁴The term *unconditional security* is borrowed from classical cryptography and reflects the notion that no conditions are imposed on the resources available to the adversary, or *eavesdropper*^[43].

ment^[2] [47–49]. Nevertheless, extending QKD protocols across large-scale quantum channels is highly challenging, since key rates decrease exponentially with distance. Direct transmission of photonic states in optical fibres suffers from loss of approximately 0.15 dB/km at telecommunication wavelengths^[50], thus limiting QKD channels to a few hundred kilometres. This prohibitive scaling can be overcome with the use of *quantum repeaters* which were first proposed by Briegel *et al.* in 1998^[21].

Quantum repeaters

The key principle underpinning the quantum repeater approach is the creation of long-range entanglement via a process known as *entanglement swapping*^[21] [51]. The quantum channel of length L is subdivided into N short elementary segments which are connected via $N-1$ intermediate nodes. As a first step, adjacent segments are independently entangled, yielding pairwise entanglement between segments 1 and 2, 3 and 4 etc. By performing a local entanglement operation on segments 2 and 3, it is possible to establish entanglement between segments 1 and 4 (see Figure 1.2). This process of entanglement swapping is then repeated in a nested fashion⁵, until the entangled state has been extended across the entire distance L .




Figure 1.2: DLCZ-type entanglement swapping operation based on a joint measurement between segments 2 and 3. Detection of a single photon heralds the creation of entanglement between remote segments 1 and 4. Dotted lines indicate existing entanglement.

⁵The number of necessary entanglement swapping operations is given by the *nesting level* n , where $N = 2^n$ [52].

The principal requirements for a quantum repeater protocol are the ability to herald successful swapping operations and subsequent storage of entanglement in quantum memories^[53–55]. In addition, one needs to be able to swap entanglement between these memories by performing local joint measurements on the photons retrieved from each memory. Provided that these measurements erase all information about the origin of each photon (known as *which-way* information^{[56] [57]}), the memories are projected onto an entangled state.

A pioneering approach for realising a quantum repeater was proposed in 2001 by Duan, Lukin, Cirac and Zoller^[58]. Combining atomic ensemble memories, linear optics and single-photon detection, the DLCZ protocol has since encouraged a profusion of experimental demonstrations^[59–64]. One of the key insights was that the communication efficiency, or the time needed to transmit an entangled state, could be made to scale polynomially with channel length, rather than exponentially.

The realisation of quantum repeater architectures is a critical step in the path towards scalable quantum networks of interconnected nodes^{[3] [65]}, enabling both local processing and long-distance distribution of quantum information^{[66] [67]}.

1.2.3 Quantum computing

Perhaps one of the most radically profound, yet elusive quantum information technologies is quantum computing^{[4] [68] [69]}. It holds the potential for exponential speed-up of certain tasks which are insurmountably difficult within the limits of classical computation. These tasks include factorisation^[70,71], search algorithms^{[72] [73]} and the simulation of the dynamics of complex quantum systems^[74].

A variety of candidate systems have been investigated for implementing QIP protocols, including trapped ions^[75–78], superconducting qubits^{[79] [80]} and solid-state ensembles, for

example rare-earth ions in crystalline hosts^[81]. Quantum optics in particular has proved to be a fertile testing ground for fundamental concepts in quantum information science since the early 1970s^{[82][83]}. While single photons have emerged as highly attractive information carriers^{[84][85]}, photonic QIP has until recently been beset by the challenge of achieving optical nonlinearities sufficiently large to perform quantum logic gate operations^[86].

However, in 2001 Knill, Laflamme and Milburn revolutionised the field of optical quantum computing by demonstrating that efficient quantum computation is possible using only single-photon sources, linear optical networks and photon-detectors^[87]. The KLM scheme overcomes the scalability problem of probabilistic two-qubit gates in linear optical quantum computing (LOQC) by harnessing quantum teleportation^{[88][89]}. This involves teleporting two qubits onto the output of a non-deterministic gate, provided that the gate operation has succeeded. In order to preserve the prepared qubits until a successful gate has been heralded, the qubits can be taken *off-line*^[90] and, for instance, interfaced with quantum memories.

More generally, quantum memories can be used to synchronise probabilistic logic gates and facilitate feed-forward operation of downstream processing steps, conditional on the success of prior qubit operations^{[91][92]}. The advantage of integrating quantum memories into scalable LOQC schemes is discussed in more detail in reference^[93].

1.3 Figures of merit

Before reviewing some of the most widely studied quantum memory schemes, we will first examine the salient metrics with which to assess and distinguish these protocols. In addition to the figures of merit listed below, technical simplicity, robustness and ease of implementation,

in particular room-temperature operation, are advantageous for integration into real-world QIP architectures.

1.3.1 Fidelity

The fidelity $F(\rho, \sigma)$ is a measure of the distance^[3] between two states with density matrices ρ and σ , defined as⁶

$$F(\rho, \sigma) = \left(\text{tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2. \quad (1.2)$$

It is straightforward to show that if one of the states is in a pure state $|\psi\rangle$, the fidelity $F(|\psi\rangle, \sigma)$ simply reduces to the overlap between $|\psi\rangle$ and σ , i.e. $F(|\psi\rangle, \sigma) = \langle \psi | \sigma | \psi \rangle$.

In the context of quantum memories, the fidelity $F(\rho_{\text{in}}, \rho_{\text{out}})$ can be used to evaluate the amount of overlap between the original input state ρ_{in} and the state ρ_{out} retrieved from the memory.

As a point of reference, the memory fidelity can be compared to the maximum fidelity achievable with a classical *measure-and-prepare* strategy^{[95][96]} which aims to reproduce a quantum state based on the result of a measurement. For a single-qubit measurement, the fidelity is bounded by $F = 2/3$, as demonstrated originally for the case of estimating the polarisation state of a photon from a projective measurement along one basis state^[97]. For an ensemble of N particles, the maximum fidelity scales as $(N + 1)/(N + 2)$ which tends towards unity as $N \rightarrow \infty$, i.e. the state can be determined exactly for an infinitely large ensemble. The fidelity boundary can be used to determine whether the retrieved state is more closely matched to the input state than any copy obtained through classical measurements. It is important to bear in mind that the fidelity is in fact a state-dependent

⁶There exists a different convention which defines the fidelity as the square root of equation (1.2). Here, we will use the definition given in the original proposal^[94].

measure^[98]. For instance, if the memory is characterised with coherent input states^{[99][100]}, the Poissonian photon number distribution has to be taken into account when determining the appropriate fidelity boundary⁷.

In Chapter 4, we will make use of a related figure of merit, the *process fidelity*^[102]. While the state fidelity, as defined above, quantifies the distance between a given input state and the corresponding output state, the process fidelity compares the measured process matrix of the memory interaction with the ideal process matrix expected if the memory were simply the identity operation. It can therefore be understood as a distance measure between the real and the ideal memory process, i.e. a measure which determines the ability of the memory to store an arbitrary state faithfully.

1.3.2 Memory efficiency

The memory efficiency η is defined as the ratio of the retrieved energy to the incident signal energy, the former of which may be corrected for potential noise contributions present during read-out of the memory. This definition of the efficiency is typically applied when coherent states are used as input signals. For probabilistic single-photon states, one can define the efficiency as the probability of retrieving a photon from the memory, conditional on the presence of a non-vacuum input state^[101]. Note that the memory efficiency generally refers to the total efficiency which is equivalent to the product of the individual efficiencies of the storage and retrieval processes.

The efficiency is one of the most commonly used figures of merit when characterising a quantum memory; however, it is insufficient when knowledge of the *closeness* between the input and output states is required. Fidelity measurements may therefore prove useful in order

⁷As shown in reference^[101] for the case of coherent states with a Poissonian distribution $p(\mu, N)$, the fidelity boundary depends on the mean photon number μ according to $F_{\text{coh}}(\mu) = \sum_{N \geq 1} \frac{N+1}{N+2} \frac{p(\mu, N)}{1-p(\mu, 0)}$.

to gain additional, albeit state-dependent information about the memory performance⁸.

Quantum repeater protocols require exceptionally high memory efficiencies ($> 90\%$) to achieve sufficient rates of entanglement swapping^[52]. However, other applications have less stringent requirements on memory efficiency and can tolerate loss more easily. As discussed in Section 1.2.1, temporal multiplexing tasks can be accomplished with significantly less efficient memories. If error-correcting algorithms are employed, loss tolerance can be improved to allow for efficiencies of only 50% ^[18], as demonstrated for one-way quantum computation with cluster states^[105–107].

1.3.3 Storage time

The maximum storage time possible with a quantum memory is limited by the timescale of decoherence and loss processes which degrade the stored atomic excitation. The memory lifetime is most commonly determined from the $1/e$ decay of the efficiency^[108]. For quantum repeater applications the lifetime should be on the order of several milliseconds^[43], and ideally be comparable to the time necessary for total entanglement distribution across the entire quantum channel^[52]. In reference^[109] it was demonstrated that the repeater rate declined as $\exp\left(-\sqrt{L/cT_s}\right)$, where L is the total channel length and T_s is the storage time in the limit $T_s \ll L/c$. However, it has been suggested that the use of multiplexing strategies could significantly reduce the memory lifetime requirement^[110].

In the case of synchronisation protocols, the memory lifetime should span several repeti-

⁸The effect of non-unit memory efficiencies on the fidelity is itself state-dependent, as discussed in reference^[103]. For example, this can be visualised by considering the phase-space diagram for coherent states; the amount of overlap between a coherent state (the input) and a second coherent state (the output), whose amplitude is reduced by some fixed fraction relative to the input state (representing memory loss), depends on the mean photon number of the initial state^[104]. In the case of single photons, fidelity and efficiency can be decoupled via post-selection, i.e. fidelity measurements are made to be conditional on detecting a photon. In addition, it has been shown that non-unit efficiencies increase the fidelity boundary set by the classical measure-and-prepare strategy^[101].

tion cycles of the probabilistic system in order to achieve a significant scaling advantage over unsynchronised protocols. For many memory applications, including temporal multiplexing, the ability to retrieve stored information on-demand is a key requirement^[21]. Active control of the storage time, as opposed to pre-determined temporal delays, is therefore an important distinguishing feature of quantum memories.

1.3.4 Bandwidth

The acceptance bandwidth δ of a quantum memory determines the minimum temporal width of the pulses that can be stored. While bandwidth requirements differ between applications, broadband storage capabilities are desirable when interfacing memories with single photon sources based on parametric downconversion whose bandwidths are typically on the order of THz^[111]. Alternatively, the bandwidth of single photon sources can be reduced to < 10 MHz using optical cavities^[112–114] in order to render these compatible with narrow-band memories.

1.3.5 Time-bandwidth product

The time-bandwidth product $B = T_s \times \delta$ is a particularly important figure of merit for temporal multiplexing applications. Defined as the ratio of storage time to pulse duration, it provides a measure of the maximum number of temporal bins available within the memory lifetime. B is therefore directly proportional to the number of probabilistic processing steps that can be attempted for synchronisation purposes. As discussed in Section 1.2.1, time-bandwidth-products on the order of ~ 1000 strongly enhance the scalability of photonic networks.

1.3.6 Multi-mode capacity

The multimode capacity describes the number of modes – whether spatial^{[115][116]}, spectral^{[34][54][117]} or temporal^{[118][119]} – which can be simultaneously stored in a quantum memory. Quantum repeater protocols greatly benefit from a large multimode capacity, as the ability to store N photons at the same time gives rise to an N -fold enhancement of the rate at which entanglement is created between distant nodes^[110]. More generally, multimode capacity in any degree of freedom can enable multiplexing strategies which, as discussed in Section 1.2.1, are pivotal for the scalability of QIP technologies^[120].

1.4 Review of quantum memories

The burgeoning development of photonic QIP protocols has led to extensive research into a wide variety of quantum memory schemes. Implemented in a range of different physical systems and media, these schemes tend to display somewhat obverse benefits and challenges. While the categorisation of memories is at times fluid, one can group the most common quantum memory approaches into *adiabatic* and *photon-echo* memories. The following overview is by no means exhaustive and will be restricted to ensemble-based memories. These have the advantage that the probability that an incident light field fails to interact with the atoms decreases exponentially with the number of atoms in the ensemble^[121]. Moreover, the practical realisation of ensemble memories tends to be less technically challenging.

However, despite the experimental complexities, quantum memory systems based on single emitters have been the subject of pioneering work over the last few years. Some of the earliest quantum memory proposals^[66] were inspired by the strong light-matter coupling that is possible by placing a single atom inside an optical cavity^[122]. More recently, po-

larisation storage of weak coherent states at the single-photon level has been demonstrated with single rubidium atoms trapped inside high-finesse optical cavities^[101]. Following on from this work, two single-atom memories embedded in spatially separate traps were entangled^[123], thereby realising an elementary quantum network. Impressive studies have also been conducted using systems of trapped ions^[124–126] and single spins in solid-state implementations, including nitrogen-vacancy centres in diamond^[127–130] and phosphorus donors in silicon crystals^{[131][132]}.

1.4.1 Optical delay lines

The most basic method of delaying an optical signal is to increase its path length using an optical fibre^{[133][134]}, a free-space delay line^[135] or a high-finesse cavity^{[136][103]}. Fibre delays typically lack tunability⁹ and are limited by high transmission losses of approximately 50% over a propagation distance of 15 km^[139], corresponding to a delay of approximately 72 μ s. A more advanced version of a delay line based on a free-space optical loop was proposed in reference^[140]. A proof-of-principle experiment demonstrated that single-photon polarisation qubits propagating inside the loop could be released after any number of round-trips via high-speed electro-optic switching^[32].

The use of high-Q optical cavities can be viewed as an extension of this approach. After cycling back and forth between the cavity boundaries, the signal is retrieved using electro-optical or non-linear methods^[136]. Tanabe *et al.* demonstrated in 2007 that the Q-factor of photonic crystal nanocavities can be dynamically tuned to provide controllable delays^[141].

Light storage in optical cavities generally suffers from a trade-off between short cycle periods and long storage times^[142]. Furthermore, intracavity losses severely limit the available

⁹Stimulated Brillouin scattering can be used to overcome this problem and allow tunable delays in optical fibres^{[137][138]}.

time-bandwidth product^[143]. Optical delay lines and cavities are therefore not compatible with the particular requirements of long-distance quantum communication, although they may prove suitable for on-demand single-photon sources.

1.4.2 Adiabatic memories

The inherent limitations of optical delay lines can be overcome by exploiting the coherent interaction between optical fields and atomic systems.

Adiabatic quantum memories take advantage of a particular type of atomic energy-level structure; the Λ -level configuration comprises two ground states coupled to a single excited state. Direct transitions between the ground states are dipole-forbidden due to angular momentum selection rules. Storage of a signal field is mediated via a strong control field which dynamically modulates the coupling of the signal to the atomic system. A key property of adiabatic memories is that the excited state is never populated during the memory interaction. As a result, the stored signal cannot be lost to spontaneous emission.

Below we will consider two limiting cases of adiabatic memory; in the case of *electromagnetically induced transparency*, the optical fields are tuned close to resonance with the atomic transitions, whereas the *Raman* interaction describes the far-off-resonant limit.

Electromagnetically Induced Transparency

Electromagnetically induced transparency (EIT) is a non-linear optical phenomenon by which an optically thick medium is rendered transparent^[144]. As pictured in Figure 1.3 (a), the control field mixes states $|2\rangle$ and $|3\rangle$, producing the dressed states $|\pm\rangle = (|2\rangle \pm |3\rangle) / \sqrt{2}$ (Figure 1.3 (b)). As a result, the absorption spectrum of a probe signal tuned to the

$|1\rangle \leftrightarrow |2\rangle$ transition exhibits a double-resonance feature. In fact a full analysis of the system's linear susceptibility χ reveals that the absorption of the probe field vanishes completely^[145], giving rise to a transparency window between the two resonance peaks, as shown in Figure 1.3 (c). This result is due to destructive interference between the probability amplitudes associated with different transition pathways which are supported by the Λ -level structure, e.g. $|1\rangle \rightarrow |2\rangle$ and $|1\rangle \rightarrow |2\rangle \rightarrow |3\rangle \rightarrow |2\rangle$ ^[53]. The width of the transparency window is thereby determined by the Rabi frequency of the control field.

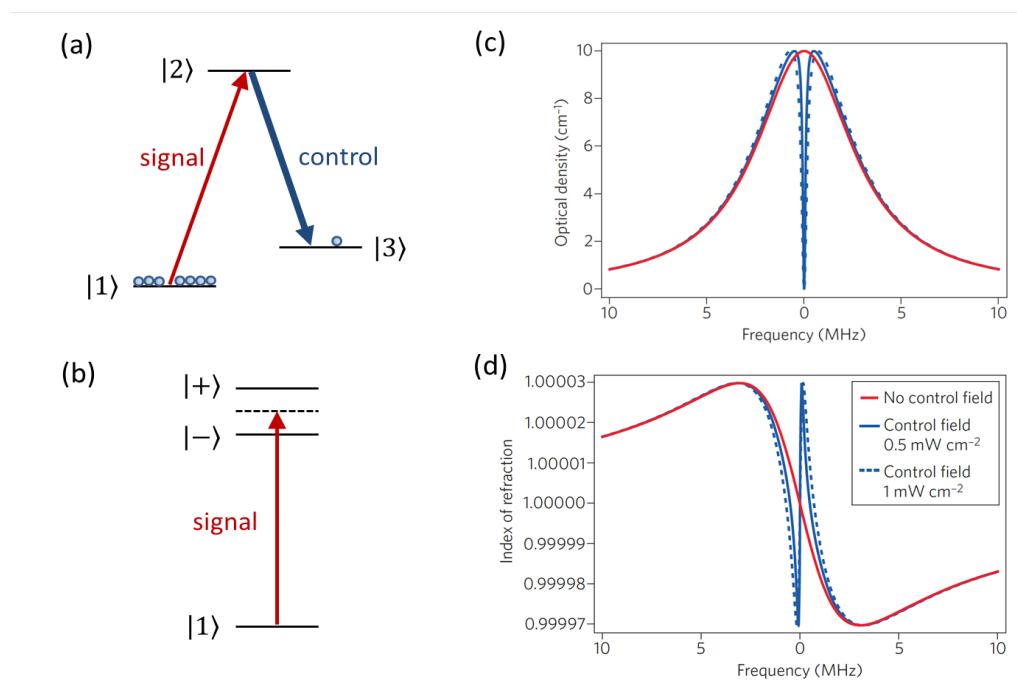


Figure 1.3: (a) EIT protocol: The signal and control fields interact with an atomic Λ -level system. (b) The control field couples states $|2\rangle$ and $|3\rangle$, thus producing the dressed states $|+\rangle$ and $|-\rangle$. The signal field is tuned in between these two dressed states. (c) Imaginary part of the linear susceptibility as a function of detuning from the atomic resonance frequency with the control field off (red line) and on (blue line), showing the transparency window at the centre of the double resonance feature. (d) Real part of the susceptibility (same colour coding), showing large positive dispersion. Images (c) and (d) are taken from reference^[142].

According to the Kramers-Kronig relations^[146], the reduction in absorption is accompanied by large positive dispersion, as shown in Figure 1.3 (d). A signal field tuned to the

centre of the transparency window will therefore experience a sharp reduction in its group velocity $v_g = d\omega/dk = c/(n + \omega dn/d\omega)$, where $n = \sqrt{1 + \text{Re}(\chi)}$ ^[147]. As a result, the pulse is spatially compressed until it completely fits inside the atomic ensemble.

In order to store the signal in the atomic medium, the control field, and hence the splitting of the dressed states, is adiabatically reduced to zero at which point the transparency window collapses and the dispersion of the probe diverges ($v_g = 0$). The quantum state of the optical field is thereby reversibly mapped onto a collective ground state coherence of the ensemble of M atoms. This can be expressed as

$$|\Psi\rangle = \sum_j \psi_j e^{i\Delta k z_j} |g_1 \dots s_j \dots g_M\rangle, \quad (1.3)$$

where z_j is the position of the j^{th} atom and Δk is the difference between the wavevectors of the control and signal fields^[142]. The ground and storage states are denoted by g and s , respectively. The amplitudes ψ_j quantify the interaction strength of the light field with the j^{th} atom within the ensemble. Upon retrieval, the signal field is re-accelerated by increasing the control field until the signal emerges from the ensemble^[53].

EIT memories were originally proposed by Fleischhauer and Lukin in 2000^[148] and have since been implemented in a wide variety of storage media, including both cold and warm atomic ensembles^[149–151] and solid state systems^{[152] [153]}.

Record coherence times of several seconds were demonstrated in rare-earth doped solids^{[154] [155]} and remarkably high efficiencies of 78% have been achieved in cold atomic ensembles^[156]. Furthermore, EIT memory for non-classical states of light has been shown with single photons generated using DLCZ methods^{[157] [158]}. Storage and retrieval of squeezed vacuum states has also been observed^[159–161]. In an impressive experiment conducted in

2008, Choi et al. demonstrated for the first time the reversible mapping between an entangled state of light and two separate quantum memories^[162].

Both theoretical^{[163][164]} and experimental^{[165][166]} studies have shown that EIT storage at the single-photon level suffers from background noise in the retrieved signal. This noise is thought to be due to residual population in the storage state $|3\rangle$ which gives rise to scattered thermal photons in the presence of the control field.

One of the requirements of EIT storage is that the spectrum of the signal pulse should not be wider than the transparency window in order to avoid partial absorption. This width is set by the available control field intensity which imposes a considerable constraint on the signal bandwidth. EIT memories are therefore typically operated in the MHz regime^[142] which means that direct interfacing with single-photon sources is a significant challenge.

Raman memory

Raman scattering describes the inelastic scattering of light from an atom or molecule. Following theoretical predictions in the early 1920s^[167], C. V. Raman – together with his collaborator K. S. Krishnan – made their eponymous discovery^[168] in 1928¹⁰. Raman scattering would prove its pervasive relevance for many areas of physics and chemistry throughout the last century, gaining particular momentum with the advent of the laser in the 1960s. Ranging from spectroscopy and microscopy, over imaging and sensing to quantum optics, the Raman effect has shown remarkable scope for diverse applications.

As a two-photon process, Raman scattering is second-order in the electric dipole interaction^[170]. If the frequency of the scattered photon is lower than the frequency of the absorbed

¹⁰In an extraordinary example of multiple discovery, two Soviet physicists – Grigory Landsberg and Leonid Mandelstam – observed the same inelastic scattering effect^[169], exactly a week before Raman reported to have made his first observation. In recognition of this, Russian scientific literature usually refers to the phenomenon as *combinatorial scattering*.

photon, the process is known as *Stokes scattering*. In the opposite case, an *anti-Stokes* photon of higher energy is emitted.

In contrast to EIT, the Raman memory interaction couples signal and control fields which are both far detuned from the excited state $|2\rangle$. Prior to storage, the atomic ensemble is prepared in the ground state $|1\rangle$ via optical pumping. During the two-photon Raman transition the signal field is absorbed into a virtual resonance created by the control field and subsequently transferred to the storage state $|3\rangle$ via Raman scattering (see Figure 1.4). The two-photon resonance condition requires that the signal and control frequencies differ by an amount equal to the ground-state splitting of the atomic Λ -system^[171–173].




Figure 1.4: The Raman memory interaction for (a) storage and (b) retrieval of a signal field with frequency ω_s , mediated via a strong control field (ω_c), both of which are detuned by Δ from the excited state $|2\rangle$.

As with other ensemble-based memory schemes, the signal field is mapped onto a collective, delocalised atomic excitation, or *spin wave*, described by equation (1.3). Upon retrieval, the interaction is reversed, i.e. re-application of the strong control pulse converts the spin wave into an optical output signal. The spectral width of the virtual state is thereby set by the bandwidth of the control pulse and only limited by the ground-state splitting of the atomic system. In the case of alkali atoms this splitting is on the order of GHz. Unlike protocols relying on EIT, the far off-resonant Raman memory is therefore capable of storing

broadband – and temporally short – photons^[174]. Assuming storage times on the order of several hundred microseconds in hot alkali vapours^[175], the time-bandwidth product of Raman memories could potentially be as high as $B = 1 \times 10^5$. Even with currently achievable storage times of a few microseconds¹¹, $B \sim O(10^3)$ is sufficiently high to be useful for temporal multiplexing applications.

A further advantage of the Raman memory is its insensitivity to inhomogeneous broadening, as the signal and control fields are detuned well beyond the Doppler-broadened linewidth of the excited state. This feature makes the Raman memory ideal for room-temperature operation which greatly simplifies practical implementations.

Finally, the far-off-resonant nature of the Raman interaction ensures that any portion of the signal not successfully stored in the memory is transmitted through the ensemble without being resonantly absorbed. As a result, the Raman memory behaves like a beam splitter with a reflectivity determined by the storage efficiency. This will be examined more rigorously in Section 2.2.4.

The Raman memory was first realised in warm caesium vapour using broadband pulses with a bandwidth of 1.2 GHz^[174]. A detailed overview of subsequent experiments will be given in Section 3.2. The Raman memory has also been demonstrated with THz-bandwidth pulses in hydrogen gas^[176] and in the optical phonon modes of bulk diamond^[177], although storage times are limited to 1 ns and 3.5 ps, respectively. The latter implementation has most recently been used to show room-temperature storage and retrieval of ultrafast, THz-bandwidth heralded single photons^[178]. However, the rapid decay of optical phonons limits the memory lifetime to a ps-timescale which is prohibitively short for many applications.

¹¹At present, the storage time is limited by magnetic dephasing of the spin coherence.

1.4.3 Photon-echo memory

As an alternative approach to light storage, a class of quantum memories has evolved around the concept of *photon echoes*. The phenomenon of spin rephasing was first discovered in nuclear magnetic resonance by Erwin Hahn in 1950^[179]. The corresponding effect at optical frequencies was initially demonstrated in two-pulse photon echo experiments^[180]. An ensemble of two-level atoms is prepared in a coherent superposition of ground and excited states using a $\pi/2$ -pulse. If the transition is inhomogeneously broadened, the initially aligned dipoles will begin to precess at different rates, thereby dephasing the atomic coherence and trapping the excitation inside the ensemble^[142]. Application of a π -pulse inverts the atomic populations so that the collective coherence will rephase after a time equal to the dephasing time. Once all the dipoles have been realigned, the original signal is re-emitted as a photon echo.

However, two-pulse photon echoes are not suitable as quantum memories, since the optical inversion pulse used to induce rephasing can lead to spontaneous emission and amplification¹² [184] [185]. This problem can be circumvented by employing a non-optical method of reversing the dephasing process of the atomic coherence via artificial control of the inhomogeneous broadening. Several different protocols have been proposed to this end, including controlled reversible inhomogeneous broadening, gradient echo memory and atomic frequency combs.

Controlled Reversible Inhomogeneous Broadening and Gradient Echo Memory

Controlled reversible inhomogeneous broadening (CRIB) was first proposed by Moiseev *et al.* in 2001^[186]. While the original proposal suggested a Doppler-broadened atomic vapour

¹²It has been shown that the noise generated by two-pulse photon echoes can be made useful in DLCZ-type quantum repeater schemes^[181–183].

as the storage medium, most subsequent implementations have been made in solid-state systems^{[117] [187–189]}. By applying a spatially varying electric or magnetic field, the $|1\rangle \leftrightarrow |2\rangle$ transition is inhomogeneously broadened across the entire ensemble, since atoms at different locations exhibit different detunings^[190]. After absorption of a signal pulse resonant with the broadened transition, the collective atomic excitation can be expressed as

$$|\Psi\rangle = \sum_j \psi_j e^{-i\delta_j t} e^{ikz_j} |g_1 \dots e_j \dots g_M\rangle, \quad (1.4)$$

where δ_j is the detuning of the transition of atom j , e is the excited state and k denotes the wavevector of the signal field^[121].

As can be seen from equation (1.4), the initial phase alignment of the atomic dipoles decays rapidly due to the different values of δ_j . By changing the polarity of the external field gradient, the frequency detunings are reversed, i.e. $\delta_j \rightarrow -\delta_j$. As a result, the dipoles rephase and the signal is re-emitted in a time-reversal of the absorption process.




Figure 1.5: CRIB protocol with spin wave storage: (a) The signal field is resonantly absorbed into the inhomogeneously broadened excited state and a strong control field transfers the excitation to the storage state $|3\rangle$. (c) Retrieval: After re-application of the control field, the polarity of the external field gradient is reversed and the stored signal is released as a photon echo.

With the addition of a dipole-decoupled storage state (labelled $|3\rangle$ in Figure 1.5), the optical excitation can be mapped to a spin-wave coherence by applying a control pulse, once

the signal field has been absorbed. Retrieval of the stored excitation is achieved by first re-applying the control field, before reversing the field gradient.

The first experimental realisation of CRIB in a europium-doped Y_2SiO_5 crystal was reported by Alexander *et al.* in 2006^[188]. Solid-state systems doped with rare-earth ions are a particularly attractive choice for memory implementations, as they offer exceptionally long coherence times^{[81][191]}. This is due to the unusual electronic structure of rare-earth atoms; the optically active and tightly-bound valence $4f^N$ electrons are shielded from the environment by the outer $5s^2$ and $5p^6$ shells^[185].

However, the difficulty of achieving high optical depths with a single absorption line has meant that memory efficiencies for CRIB protocols are typically very low^{[192][193]}.

A variant of CRIB is the gradient echo memory (GEM) which was proposed by Hetet *et al.* in 2008^[189]. While the CRIB scheme described above employs transverse broadening of the ensemble, GEM involves a monotonic field gradient applied parallel to the direction of signal propagation. This was motivated by the discovery that a longitudinally broadened ensemble could prevent re-absorption of the retrieved photons so that – in principle – perfect read-out efficiencies could be achieved. In addition, GEM maps each frequency component of the stored signal pulse to a distinct location along the ensemble axis, resulting in a frequency-encoding memory as shown in Figure 1.6.

Following initial experiments in two-level rare-earth doped solids^[194], a three-level version known as Λ -GEM was realised in warm rubidium vapour^{[104][195]} with a total memory efficiency of 87%, the highest value recorded to date^[196]. Furthermore, storage of coherent states at the single photon level has been demonstrated with a fidelity of 98%^[98]. However, true quantum storage of heralded single photons remains challenging due to the typically

low memory bandwidths on the order of 1 MHz¹³.




Figure 1.6: GEM protocol: (a) An external field gradient is applied longitudinally along the memory axis. (b) Reversal of the field gradient at a time $t = T$ leads to re-emission of the signal at $t = 2T$.

Atomic Frequency Comb

The atomic frequency comb (AFC) memory is largely based on the CRIB scheme, however it requires the preparation of an ensemble with a large number of equally spaced absorption lines. Afzelius *et al.* suggested that this comb-like absorption structure could be realised in rare-earth-doped crystals which naturally exhibit a wide range of resonant frequencies due to spatially varying frequency shifts within the host matrix^[120]. The frequency comb can be prepared via optical pumping in order to shelve atoms with resonant frequencies in between the comb teeth, thereby removing these atoms from the memory process.

While the AFC storage process is similar to the CRIB protocol, the retrieval of the photon echo occurs periodically without the need to reverse the detunings of the atoms. This is due to the discrete spectral structure of the frequency comb which results in periodic rephasings of the atomic spins with a periodicity of $2\pi/\Delta$, where Δ is the frequency spacing of the comb teeth. The timing of the photon echo in a two-level system is therefore fixed by the comb structure, unless the excitation can be transferred to a third storage state.

¹³The bandwidth of GEM protocols is set by the product of the applied field gradient and the length of the ensemble. The bandwidth is therefore easily tunable, although it should be noted that for a given optical depth, a trade-off exists between the bandwidth and the memory efficiency^[98].




Figure 1.7: Three-level AFC protocol: (a) A signal field with a bandwidth $\delta \gg \Delta$ is absorbed by the absorption lines of the frequency comb and then mapped to the storage state $|3\rangle$. (b) Due to the discrete spacing of comb teeth, photon echoes are released periodically at multiples of $2\pi/\Delta$.

The first AFC memory for weak coherent states was demonstrated in neodymium-doped YVO_4 ^[197] and high efficiencies of up to 56% have since been achieved^[198].

One of the key advantages of the AFC scheme is its high multimode capacity which scales with the number of absorption peaks spanned by the signal bandwidth^[120]. Simultaneous storage of 1060 pulses as well as bandwidths of up to 5 GHz have been reported^{[199][200]}. Moreover, single-photon-level memory operation at telecommunication wavelengths^[193], storage of entangled photons^[201] and heralded entanglement between two separate AFC memories^[202] have been demonstrated.

These experiments have so far been restricted to two-level AFC protocols without spin-wave storage. In order to move beyond the limitations of pre-determined storage times and enable on-demand retrieval, recent work has focused on demonstrating three-level AFC memories^{[203][204]}. As illustrated in Figure 1.7, the initial excitation is reversibly mapped onto a spin wave by applying a control pulse whose bandwidth spans the AFC width. After a variable delay T , a second control pulse then returns the excitation to the frequency comb, resulting in a photon echo at $T + 2\pi/\Delta$. While current memory efficiencies are still very low ($\sim 6\%$), there have been some encouraging improvements recently^[205–207].

1.5 Thesis outline

This thesis is concerned with two experimental realisations of the far-off-resonant Raman memory, first in free space and then – in response to the results obtained – inside an optical cavity. Below I will give an overview of the structure of my thesis and summarise the contents of each chapter.

A theoretical description of the Raman memory will be given in the following chapter, beginning with a derivation of the optical Bloch equations which govern the dynamics of the atomic system. Based on these, we can develop a set of coupled equations of motion which describe the memory interaction in the adiabatic regime. The optical and spin wave modes coupled by this interaction are given by a linear mapping from which the optimal memory parameters can be determined.

The remainder of this thesis is divided into two parts which appertain to the free-space and cavity implementations of the Raman memory. The chapters in both parts follow a chronological order, as the work presented in each builds upon and is motivated by preceding results. Common to all is the overarching aim to realise a Raman memory capable of availing quantum information tasks.

The experimental realisation of the Raman memory is detailed in Chapter 3, including the memory medium, the optical set-up as well as the operating procedure. This is followed by a summary of previous work, in particular the first demonstration of broadband coherent-state storage, experiments at the single-photon level and the application of Raman memories as dynamically configurable beam splitter networks.

Chapter 4 describes the high-fidelity storage of coherent polarisation states in two spatially separate memory ensembles contained inside a dual-rail polarisation interferometer.

Preservation of polarisation coherence is demonstrated by performing full process tomography on the memory interaction.

Having so far examined the Raman memory with classical, coherent states of light, we will next consider the storage of broadband heralded single photons generated from spontaneous parametric downconversion. Chapter 5 describes how a non-linear waveguide source of heralded single photons is interfaced with the Raman memory to enable memory operation conditional on detection of a herald photon. The photon number statistics of the retrieved field are determined from second-order autocorrelation measurements and compared to those obtained for coherent-state signal fields. However, the preservation of non-classical photon statistics is impaired by the presence of four-wave mixing noise. The statistics of this noise contribution can be modelled by taking into account the full coherent interaction of the off-resonant optical fields and the spin wave excitation. An experimental investigation of the noise process is then presented to further consolidate our understanding of the underlying mechanism. The last part of Chapter 5 discusses the feasibility of different strategies for reducing this source of noise with the aim of allowing Raman memories to be used in future quantum information applications.

This analysis encourages us to depart from a free-space implementation and instead consider memory operation inside a low-finesse optical cavity designed to suppress four-wave mixing noise. Chapter 6 introduces the necessary formalism of input-output theory and lays out the dynamical equations describing the cavity Raman memory interaction. This is followed by a discussion of the cavity-induced enhancement of the coupling strength as well as the heightened mode selectivity of the system. The latter half of Chapter 6 details the design specifics of the cavity implementation. This entails geometrical requirements, parameter optimisation and stability considerations.

Finally, the experimental realisation of the cavity Raman memory is presented in Chapter 7, including a description of all optical components, as well as the cavity stabilisation method and memory operation. Storage and retrieval of weak coherent states is demonstrated for the first time and suppression of four-wave mixing noise is characterised with the present cavity design.

Chapter 2

Theoretical overview of the Raman memory

“The miracle of the appropriateness of the language of mathematics for the formulation of the laws of physics is a wonderful gift which we neither understand nor deserve.”

Eugene Wigner

The theory of the Raman memory interaction was first laid out by Kozhokin *et al.*^[173] and subsequently extended by Gorshkov^[208] and Nunn in 2007^[170]. The following overview, which is based on^[170], begins with a description of the Hamiltonian for the interaction between a weak signal field and a system of Λ -level atoms in the presence of a strong control beam.

The Λ -level structure consists of two dipole-decoupled ground states, labelled $|1\rangle$ and $|3\rangle$, which are connected via a single excited state $|2\rangle$. This simple configuration is at the very core of a rich variety of non-linear optical phenomena, including electromagnetically-induced transparency^{[144][145]}, coherent population trapping^{[209][210]} and stimulated Raman adiabatic passage (STIRAP)^{[211][212]}, to name but a few. STIRAP describes a method of

coherently transferring population between the ground states $|1\rangle$ and $|3\rangle$ and the physics of this process is remarkably pertinent to both Raman and EIT memories^[213].

Once we have found an expression for the light-matter Hamiltonian, we can establish the dynamical equations which govern the atomic evolution. The propagation of the signal field through the atomic ensemble can be determined from Maxwell's equations. By defining continuous atomic operators for the atomic polarisation and spin wave, we finally arrive at the equations of motion which describe the macroscopic behaviour of the Raman memory.

2.1 Equations of motion

2.1.1 The light-matter Hamiltonian

The complete Hamiltonian of the light-matter system can be written as the sum of three separate contributions,

$$\hat{H} = \hat{H}_A + \hat{H}_L + \hat{H}_{ED}, \quad (2.1)$$

where \hat{H}_A is the Hamiltonian for the bare atom, \hat{H}_L is the free-field Hamiltonian and \hat{H}_{ED} is the electric dipole Hamiltonian.

The first of these, the Hamiltonian for the bare atom, can be diagonalised by expressing \hat{H}_A in terms of its eigenstates $|1\rangle$, $|2\rangle$ and $|3\rangle$,

$$\hat{H}_A = \sum_j \hbar\omega_j |j\rangle \langle j| = \sum_j \hbar\omega_j \hat{\sigma}_{jj}, \quad (2.2)$$

where ω_j is the resonant frequency of the eigenstate $|j\rangle$. This is known as *second quantisation*

of the atomic Hamiltonian^[214] and can be derived by applying the closure theorem^[215]

$$\sum_j |j\rangle \langle j| = 1 \quad (2.3)$$

twice to the eigenvalue equation $\hat{H}_A |i\rangle = \hbar\omega_i |i\rangle$, giving

$$\hat{H}_A = \sum_j |j\rangle \langle j| \hat{H}_A \sum_k |k\rangle \langle k|. \quad (2.4)$$

Using the orthonormality condition for the eigenstates $\langle j| \hat{H}_A |k\rangle = \hbar\omega_j \delta_{jk}$, equation (2.4) reduces to (2.2), where $\hat{\sigma}_{jk} = |j\rangle \langle k|$ is the atomic ‘flip’ operator, also known as the *transition* or *projection* operator^[216]. This nomenclature arises from the fact that application of $\hat{\sigma}_{jk}$ projects an atomic state $|k\rangle$ onto the orthonormal state $|j\rangle$, since

$$\hat{\sigma}_{jk} |l\rangle = |j\rangle \langle k|l\rangle = |j\rangle \delta_{kl} = |j\rangle \quad \text{if} \quad k = l. \quad (2.5)$$

It is straightforward to verify that the flip operators satisfy the following relations:

$$\hat{\sigma}_{ij} \hat{\sigma}_{kl} = \hat{\sigma}_{il} \delta_{jk} \quad \text{and} \quad \hat{\sigma}_{jk}^\dagger = \hat{\sigma}_{kj}. \quad (2.6)$$

The free-field, or *radiative* Hamiltonian \hat{H}_L describes excitations of the electromagnetic field and is given by

$$\hat{H}_L = \int_0^\infty d\omega \hbar\omega \hat{a}^\dagger(\omega) \hat{a}(\omega). \quad (2.7)$$

Here, the zero-point energy has been ignored and the annihilation operator $\hat{a}(\omega)$ for a photon

with frequency ω satisfies the bosonic commutation relation

$$[\hat{a}(\omega), \hat{a}^\dagger(\omega')] = \delta(\omega - \omega'). \quad (2.8)$$

The final contribution to the total Hamiltonian stems from the electric-dipole interaction described by $H_{ED} = -e\hat{\mathbf{D}} \cdot \hat{\mathbf{E}}$, where $\hat{\mathbf{E}}$ is the electric field operator and $\hat{\mathbf{D}}$ represents the dipole operator. We shall see in Section 2.1.3 that this expression is the result of the *electric dipole approximation*.

2.1.2 Optical fields

The bright control field with frequency ω_c can be expressed classically as

$$\mathbf{E}_c(t, z) = \mathbf{v}_c E_c(t, z) e^{i\omega_c(t-z/c)} + c.c., \quad (2.9)$$

where $t - z/c$ is the retarded time in a moving frame of reference and \mathbf{v}_c is the polarisation vector of the control field. For the purpose of this analysis, the slowly varying amplitude $E_c(t, z)$ is confined to one-dimensional propagation. The one-dimensional model is valid for beams of low divergence whose transverse profile can be treated as approximately constant^[47]. This greatly simplifies the following derivation¹.

Unlike the control field, the weak signal may in principle be a non-classical state which requires a quantum mechanical treatment. Expressed in terms of the annihilation operator $\hat{a}(\omega, t)$, it is given by

$$\hat{\mathbf{E}}_s(z, t) = i\mathbf{v}_s \int d\omega g(\omega) \hat{a}(\omega, t) e^{-i\omega z/c} + h.c., \quad (2.10)$$

¹An extension of this model to three dimensions has been developed in reference^[217].

where \mathbf{v}_s is the signal polarisation vector and $g(\omega) = \sqrt{\hbar\omega/4\pi\epsilon_0c\mathcal{A}}$ is the signal mode amplitude with transverse mode area \mathcal{A} . Note that we have included the time dependence of the operators in the Heisenberg picture. Assuming that the spectral width of the signal pulse is small compared to its central frequency ω_s , we can neglect the frequency dependence of the mode amplitude and re-write the signal as

$$\hat{\mathbf{E}}_s(z) = i\mathbf{v}_s g_s \hat{S}(t, z) e^{i\omega_s(t-z/c)} + h.c. \quad (2.11)$$

Here, $g_s = \sqrt{2\pi}g(\omega_s)$ and \hat{S} is the annihilation operator in the time domain which can be identified as the mode amplitude of the signal:

$$\hat{S}(t, z) = e^{-i\omega_s(t-z/c)} \times \frac{1}{\sqrt{2\pi}} \int d\omega \hat{a}(\omega, t) e^{-i\omega z/c}. \quad (2.12)$$

Equation (2.12) makes use of the slowly varying envelope (SVE) approximation, according to which the signal propagation is determined by the slow temporal evolution of S , while the exponential factor represents the rapidly oscillating carrier wave.

Furthermore, we will assume that the signal propagates along the z -axis with negligible divergence, provided that the transverse spatial profile of S is large compared to the signal wavelength. This is known as the paraxial approximation.

2.1.3 The dipole interaction

The potential energy of an atom with valence electron at \mathbf{r}_e and charge density $\sigma(\mathbf{r}) = -e\delta(\mathbf{r} - \mathbf{r}_e)$ in an electric field $\mathbf{E}(\mathbf{r})$ is given by

$$V_E = - \int d\mathbf{r} \mathbf{P}(\mathbf{r}) \cdot \mathbf{E}(\mathbf{r}), \quad (2.13)$$

where the polarisation $\mathbf{P}(\mathbf{r})$ is related to the charge density via $\sigma(\mathbf{r}) = -\nabla \cdot \mathbf{P}(\mathbf{r})$ ^[216].

We can re-write equation (2.13) in terms of a Taylor expansion of the electric field vector,

$$V_E = e \left\{ 1 + \frac{1}{2!} \mathbf{r}_e \cdot \nabla + \frac{1}{3!} (\mathbf{r}_e \cdot \nabla)^2 + \dots \right\} \mathbf{r}_e \cdot \mathbf{E}(0). \quad (2.14)$$

This expression makes explicit the contributions to the potential energy V_E from *multiple moments* of the charge distribution^[216]. In the electric dipole approximation, only the first of these terms – the *dipole* moment – is retained. This approximation is valid if the wavelength of the electromagnetic radiation is large compared to the size of the atom. It can be shown that higher-order multipole contributions to the potential energy are smaller than the electric dipole energy by factors of the fine structure constant $\alpha \ll 1$ and can therefore be dropped².

Hence, we arrive at the electric dipole interaction Hamiltonian $H_{ED} = -e\hat{\mathbf{D}} \cdot \hat{\mathbf{E}}$, introduced in Section 2.1.1. The dipole operator $\hat{\mathbf{D}}$ can be expressed in terms of its matrix elements \mathbf{d}_{jk} as

$$\hat{\mathbf{D}} = \sum_{j,k} \langle j | \hat{\mathbf{D}} | k \rangle \hat{\sigma}_{jk} = \sum_{j,k} \mathbf{d}_{jk} \hat{\sigma}_{jk}. \quad (2.15)$$

Since the electric dipole interaction has odd parity, only off-diagonal elements of the dipole operator are non-zero^[216]. We can further simplify equation (2.15) by requiring that transitions between the two ground states $|1\rangle$ and $|3\rangle$ are dipole-forbidden. The dipole operator is then given by

$$\hat{\mathbf{D}} = \mathbf{d}_{12} \hat{\sigma}_{12} + \mathbf{d}_{23} \hat{\sigma}_{23} + h.c. \quad (2.16)$$

²By retaining higher-order terms in the multipole expansion of the electric field, one can derive the more complete Power-Zienau-Woolley Hamiltonian^[216].

2.1.4 Optical Bloch equations

Having determined the light-matter Hamiltonian, we can now use the Heisenberg equation of motion to determine the temporal evolution of the atomic flip operators. We first note that these commute with the free-field Hamiltonian;

$$\partial_t \hat{\sigma}_{jk} = i[\hat{\sigma}_{jk}, \hat{H}_A + \hat{H}_L + \hat{H}_{ED}] = i[\hat{\sigma}_{jk}, \hat{H}_A + \hat{H}_{ED}] \quad (2.17)$$

Using the properties of the flip operators, (2.17) yields a set of coupled equations which describe the dynamics of the atomic populations $\hat{\sigma}_{jj}$ and atomic coherences $\hat{\sigma}_{jk}$ ($j \neq k$), i.e.

$$\begin{aligned} \partial_t \hat{\sigma}_{11} &= -i\hat{\mathbf{E}} \cdot (\mathbf{d}_{12}\hat{\sigma}_{12} - h.c.) \\ \partial_t \hat{\sigma}_{33} &= i\hat{\mathbf{E}} \cdot (\mathbf{d}_{23}\hat{\sigma}_{23} - h.c.) \\ \partial_t \hat{\sigma}_{12} &= i\omega_{21}\hat{\sigma}_{12} - i\hat{\mathbf{E}} \cdot [\mathbf{d}_{12}^*(\hat{\sigma}_{11} - \hat{\sigma}_{22}) + \mathbf{d}_{23}\hat{\sigma}_{13}] \\ \partial_t \hat{\sigma}_{13} &= i\omega_{31}\hat{\sigma}_{13} - i\hat{\mathbf{E}} \cdot [\mathbf{d}_{23}^*\hat{\sigma}_{12} - \mathbf{d}_{12}^*\hat{\sigma}_{23}] \\ \partial_t \hat{\sigma}_{23} &= i\omega_{32}\hat{\sigma}_{23} - i\hat{\mathbf{E}} \cdot [\mathbf{d}_{23}^*(\hat{\sigma}_{22} - \hat{\sigma}_{33}) - \mathbf{d}_{12}\hat{\sigma}_{13}], \end{aligned} \quad (2.18)$$

where $\omega_{jk} = \omega_j - \omega_k$ is the difference in resonance frequency between the eigenstates $|j\rangle$ and $|k\rangle$. An additional constraint is imposed by the normalisation condition $\hat{\sigma}_{11} + \hat{\sigma}_{22} + \hat{\sigma}_{33} = 1$; hence, $\partial_t \hat{\sigma}_{22} = -\partial_t(\hat{\sigma}_{11} + \hat{\sigma}_{33})$.

A number of approximations can be applied to simplify these equations. First of all, the atomic populations can be assumed to be constant in time, provided that the number of atoms far exceeds the number of photons stored in the memory³. We will therefore replace the operators $\hat{\sigma}_{jj}$ with their expectation values, i.e. $\hat{\sigma}_{11} \rightarrow 1$ and $\hat{\sigma}_{22}, \hat{\sigma}_{33} \rightarrow 0$.

The atomic coherences can be expressed in the rotating frame through the following

³Typically, the number of atoms in the ensemble is on the order of $O(10^{12})$.

transformation:

$$\tilde{\sigma}_{jk} = \hat{\sigma}_{jk} e^{i\omega_{jk}\tau}, \quad (2.19)$$

where $\tau = t - z/c$ is the retarded time. From here on, operators in the rotating frame will be marked with a tilde⁴. Inserting the expression for the electric field $\hat{\mathbf{E}}$ into the equations of motion reveals terms which oscillate rapidly at frequencies given by the sum of the optical and resonant frequencies. Using the rotating-wave approximation, these terms can be neglected if the detuning Δ of the signal and control fields from the excited state $|2\rangle$ is sufficiently small compared to the optical frequencies.

For the time being, we will also neglect any coupling of the signal field to the $|2\rangle \leftrightarrow |3\rangle$ transition as well as coupling of the control field to the $|1\rangle \leftrightarrow |2\rangle$ transition. The latter process generates spontaneous anti-Stokes scattering and a spurious contribution to the σ_{13} -coherence which is uncorrelated with the signal field. The resulting four-wave mixing process is a significant source of noise during the memory interaction and will be treated thoroughly in Chapter 5.

Finally, only terms which are linear both in the signal field \hat{S} and the coherences $\tilde{\sigma}_{jk}$ will be considered, while all second-order perturbative corrections to the dynamical equations are neglected^[170]. We therefore arrive at the following set of equations:

$$\begin{aligned} \partial_t \tilde{\sigma}_{12} &= \mathbf{d}_{12}^* \cdot \mathbf{v}_s g_s \hat{S} e^{-i\Delta\tau} - i\mathbf{d}_{23} \cdot \mathbf{v}_c E_c e^{-i\Delta\tau} \tilde{\sigma}_{13} \\ \partial_t \tilde{\sigma}_{13} &= -i\mathbf{d}_{23}^* \cdot \mathbf{v}_c^* E_c^* e^{i\Delta\tau} \tilde{\sigma}_{12} \end{aligned} \quad (2.20)$$

These equations clearly elucidate the dynamics of the memory process; the signal field \hat{S} directly excites the coherence $\tilde{\sigma}_{12}$ which is in turn coupled to the ground-state coherence

⁴We will omit the hat from our operator notation when expressed in the rotating frame.

$\tilde{\sigma}_{13}$, i.e. the spin-wave excitation between states $|1\rangle$ and $|3\rangle$, as introduced in Chapter 1.

2.1.5 Signal field propagation

From Maxwell's equations we can derive the wave equation for the propagation of the signal field \hat{E}_s through the atomic ensemble which acts as a dielectric medium in the presence of the signal:

$$\left[\nabla^2 - \frac{1}{c^2} \partial_t^2 \right] \hat{\mathbf{E}}_s = \mu_0 \partial_t^2 \hat{\mathbf{P}}_s. \quad (2.21)$$

Here, $\hat{\mathbf{P}}_s$ is the dipole moment per unit volume, or polarisation density, which acts as a source for the signal field and oscillates at the signal frequency ω_s . We will only consider this component of the total polarisation density, since we can assume that the propagation of the strong control field is unaffected by the interaction with the atomic medium. $\hat{\mathbf{P}}_s$ can be re-written in terms of the slowly varying atomic polarisation $\hat{\mathbf{P}}_s = \tilde{\mathbf{P}}_s e^{i\omega_s \tau}$.

Inserting the expression for the signal field into the wave equation and applying the paraxial and SVE approximations gives

$$\left[\nabla_{\perp}^2 + \partial_z^2 - \frac{1}{c^2} \partial_t^2 \right] \left[i\mathbf{v}_s g_s \hat{S} e^{i\omega_s(t-z/c)} \right] = \mu_0 \partial_t^2 \tilde{\mathbf{P}}_s e^{i\omega_s(t-z/c)}, \quad (2.22)$$

where, for the sake of completeness, we have included the transverse Laplacian⁵ $\nabla_{\perp}^2 = \partial_x^2 + \partial_y^2$. This expression can be simplified to

$$\left(\frac{i}{2k_s} \nabla_{\perp}^2 + \partial_z + \frac{1}{c} \partial_t \right) \hat{S} = -\frac{\mu_0 \omega_s^2}{2g_s k_s} \mathbf{v}_s^* \cdot \tilde{\mathbf{P}}_s, \quad (2.23)$$

where $k_s = \omega_s/c$ and we have again made use of the SVE approximation.

⁵In our one-dimensional propagation model the term involving the transverse Laplacian will be set to zero.

2.1.6 Raman memory equations of motion

We now define the collective polarisation operator P and the spin wave operator B by summing the contributions to the coherences σ_{jk} across the atomic ensemble:

$$\begin{aligned} P &= \frac{1}{\sqrt{n\delta V}} \sum_{\beta(\mathbf{r})} \tilde{\sigma}_{12}^{\beta} e^{i\Delta\tau} \\ B &= \frac{1}{\sqrt{n\delta V}} \sum_{\beta(\mathbf{r})} \tilde{\sigma}_{13}^{\beta}, \end{aligned} \tag{2.24}$$

where n is the number density of the atomic ensemble and the index β runs over all atoms within the small volume δV at position \mathbf{r} .

In the continuum limit $\delta V \rightarrow 0$, these operators satisfy bosonic commutation relations and can be understood as annihilation operators for the atomic polarisation and the spin wave coherence^[170], respectively:

$$\begin{aligned} [P(z, t), P^{\dagger}(z', t')] &= \delta(z - z') \\ [B(z, t), B^{\dagger}(z', t')] &= \delta(z - z'). \end{aligned} \tag{2.25}$$

The use of creation and annihilation operators in the description of the spin coherence is based on the Holstein-Primakhoff approximation^{[218] [219]}. Substituting P and B into the dynamical equations (2.20) and the propagation equation (2.23), yields the following set of equations of motion for the memory:

$$\begin{aligned} \left(\frac{i}{2k_s} \nabla_{\perp}^2 + \partial_z + \frac{1}{c} \partial_t \right) S &= -\kappa^* P \\ \partial_t P &= i\Delta P + \kappa S - i\Omega B \\ \partial_t B &= -i\Omega^* P, \end{aligned} \tag{2.26}$$

where Ω is the Rabi frequency of the control field, defined as

$$\Omega = \frac{\mathbf{d}_{23} \cdot \mathbf{v}_c}{\hbar} E_c \quad (2.27)$$

and the coupling constant κ is given by

$$\kappa = \frac{\mathbf{d}_{12}^* \cdot \mathbf{v}_s}{\hbar} \times \sqrt{\bar{n}} g_s = \frac{\mathbf{d}_{12}^* \cdot \mathbf{v}_s}{\hbar} \sqrt{\frac{\hbar \omega_s \bar{n}}{2 \epsilon_0 c}}. \quad (2.28)$$

Notice that κ incorporates the factor $\sqrt{\bar{n}}$ which accounts for the collective enhancement of the atom-field coupling in an ensemble of atoms^[53].

So far we have not considered any decay processes affecting the atomic polarisation and the spin wave coherence, such as spontaneous emission from the excited state or collisional de-phasing. These can be incorporated as follows:

$$\begin{aligned} \left(\frac{i}{2k_s} \Delta_{\perp}^2 + \partial_z + \frac{1}{c} \partial_t \right) S &= -\kappa^* P \\ \partial_t P &= -\gamma P + i\Delta P + \kappa S - i\Omega B \\ \partial_t B &= -\gamma_B B - i\Omega^* P, \end{aligned} \quad (2.29)$$

where γ is the rate of spontaneous emission and γ_B is the rate at which the spin wave decoheres.

It should be noted that the Langevin noise operators, which describe independent fluctuations of P and B while preserving bosonic commutation relations, have not been included⁶.

This is justified given that these are initially in the vacuum state so that the expectation

⁶The introduction of δ -correlated Langevin noise operators in the dynamical equations for P and B is necessary when including dissipative processes. The expectation value of normally ordered products of these operators is of the form $\langle F^\dagger(t, z) F(t', z') \rangle = 2\gamma \bar{n} \times \delta(t - t') \delta(z - z')$, where \bar{n} is the initial number of excitations. Since $\bar{n} = 0$, we can therefore neglect the Langevin noise terms when dealing with the normally ordered products $P^\dagger P$ and $B^\dagger B$.

value of terms involving normally ordered products of the Langevin operators are equal to zero^[170].

2.2 The Raman memory interaction

In order to solve the system of partial differential equations more readily we will further simplify our model by restricting the following treatment to a single dimension, thus allowing us to drop the transverse Laplacian. Re-writing the equations in terms of the retarded time $\tau = t - z/c$ and defining the complex detuning $\Gamma = \gamma - i\Delta$ gives

$$\begin{aligned}\partial_z S(z, \tau) &= -\kappa^* P(z, \tau) \\ \partial_\tau P(z, \tau) &= -\Gamma P(z, \tau) + \kappa S(z, \tau) - i\Omega(\tau) B(z, \tau) \\ \partial_\tau B(z, \tau) &= -i\Omega^*(\tau) P(z, \tau),\end{aligned}\tag{2.30}$$

where the time-dependence of the Rabi frequency $\Omega(\tau)$ accounts for the temporal profile of the control field. In the last of these equations, we have also dropped the decay term $\gamma_B B$, since decoherence of the spin wave is negligible on timescales of the storage interaction. However, spin wave decay should be taken into account when modelling the retrieval process after some storage time T_s (see Section 2.2.1). As we shall see in Section 3.2.2, this is particularly significant if T_s exceeds a few tens of ns.

2.2.1 Storage and retrieval

Having established a linear set of equations for the atomic coherences, we are now in a position to find expressions for S , P and B . This can be achieved by determining the Green's function $K(z, \tau)$ which relates the initial input signal field $S_{\text{in}}(\tau) = S(z = 0, \tau)$

and the output spin wave $B_{\text{out}}(z) = B(z, \tau \rightarrow \infty)$ via a linear transformation. $K(z, \tau)$ is the integral kernel of the linear map $S \rightarrow B$ and contains complete information about the storage interaction:

$$B_{\text{out}}(z) = \int_{-\infty}^{\infty} d\tau K(z, \tau) S_{\text{in}}(\tau). \quad (2.31)$$

A similar linear map K_{ret} connects the output field $S_{\text{out}}(\tau)$, retrieved after a storage time T_s , to the stored spin wave $B_{\text{in}}(z) = \eta_T B_{\text{out}}(z)$, where η_T represents the decay of the spin wave coherence over a time T_s . Note that in this notation $B_{\text{out}}(z)$ denotes the spin wave at the end of the read-in interaction, while $B_{\text{in}}(z)$ refers to the spin wave at the start of the read-out interaction (see Figure 2.1 below for a graphical illustration of this notation).

Hence,

$$S_{\text{out}}(\tau) = \int_0^L dz K_{\text{ret}}(z, \tau) B_{\text{in}}(z), \quad (2.32)$$

where L is the length of the ensemble. Finally, the overall memory operation of storage followed by retrieval can be described by the linear mapping

$$S_{\text{out}}(z) = \int_{-\infty}^{\infty} d\tau' K_{\text{tot}}(\tau, \tau') S_{\text{in}}(\tau'). \quad (2.33)$$

The efficiency of the storage interaction can be expressed as the ratio of the number of spin wave excitations to the number of input photons, $\eta_{\text{stor}} = N_B/N_{S,\text{in}}$, where

$$N_{S,\text{in}} = \int_{-\infty}^{\infty} d\tau S_{\text{in}}^\dagger(\tau) S_{\text{in}}(\tau) \quad \text{and} \quad N_B = \int_{-\infty}^{\infty} dz B_{\text{out}}^\dagger(z) B_{\text{out}}(z). \quad (2.34)$$

Likewise, the total memory efficiency is given by the ratio $\eta_{\text{tot}} = N_{S,\text{out}}/N_{S,\text{in}}$ where the

number of output photons $N_{S,\text{out}}$ is equal to

$$N_{S,\text{out}} = \int_{-\infty}^{\infty} d\tau S_{\text{out}}^\dagger(\tau) S_{\text{out}}(\tau). \quad (2.35)$$

The optimal storage efficiency η_{opt} as well as the optimal temporal input mode can be computed numerically by performing a singular value decomposition (SVD) of $K(z, \tau)$ ^[170]. Setting $N_S = 1$, it can be shown that η_{opt} is given by the square of the largest singular value of $K(z, \tau)$.

Initially, the atomic ensemble is prepared in one of the ground states, i.e. $\langle \sigma_{11} \rangle = 1$, $\langle \sigma_{22} \rangle = \langle \sigma_{33} \rangle = 0$, so that P and B are zero prior to the storage interaction,

$$P_{\text{in}}(z) = P(z, \tau \rightarrow -\infty) = 0 \quad (2.36)$$

$$B_{\text{in}}(z) = B(z, \tau \rightarrow -\infty) = 0.$$

Spontaneous emission from the excited state $|2\rangle$ limits the maximally achievable storage efficiency, even if unlimited control field energy were available. In order to determine this upper bound we ignore any coupling to the storage state $|3\rangle$ and simply solve the dynamical equation for P . For optimal storage we then assume perfect transfer between the atomic polarisation and the spin coherence ($P_{\text{out}} \rightarrow B_{\text{out}}$). After Fourier-transforming the variables from z -space into k -space and re-normalising the coordinates⁷ we arrive at the following map:

$$\tilde{B}_{\text{out}}(k) = \int_{-\infty}^{\infty} d\tau \tilde{K}(k, \tau_s - \tau) S_{\text{in}}(\tau), \quad (2.37)$$

where τ_s denotes the end of the storage interaction. The dimensionless Green's function is

⁷It is convenient to normalise the space and time coordinates in units of the ensemble length L and the spontaneous decay rate γ , respectively.

given by

$$\tilde{K}(k, \tau) = \frac{i\sqrt{d}}{k\sqrt{2\pi}} e^{-i(1+id/k)\tau}, \quad (2.38)$$

where κ is assumed to be real, $d = |\kappa|^2 L/\gamma$ is the resonant optical depth and L is the length of the atomic ensemble. This expression for \tilde{K} demonstrates that the optimal storage efficiency depends solely on the optical depth d which parametrises the resonant optical coupling. Experimentally, d is a measure of the fractional change in intensity of resonant monochromatic radiation passing through a partially transparent medium, i.e. $I_{\text{out}}/I_{\text{in}} = e^{-2d}$. Gorshkov *et al.* interpreted the above result by noting that the branching ratio between absorption and spontaneous emission is set by the optical depth which therefore determines the optimal storage efficiency^[208].

2.2.2 Adiabatic limit

In the adiabatic regime the detuning Δ is much larger than the Rabi frequency Ω as well as the bandwidth of the pulses. In this limit the excited state can be adiabatically eliminated by setting $\partial_t P = 0$. Solving for P and substituting the result into the remaining dynamical equations gives, in normalised units,

$$\begin{aligned} \left(\partial_z + \frac{d}{\Gamma}\right) S &= i \frac{\Omega\sqrt{d}}{\Gamma} B \\ \left(\partial_\tau + \frac{|\Omega|^2}{\Gamma}\right) B &= -i \frac{\Omega^*\sqrt{d}}{\Gamma} S. \end{aligned} \quad (2.39)$$

In the adiabatic limit of large detunings, the Green's function reduces to $K(z, \omega) = C J_0(2C\sqrt{\omega z})$, where the Raman coupling constant C is defined as

$$C = \frac{\sqrt{Wd}}{|\Gamma|} \approx \frac{\sqrt{Wd}}{\Delta}. \quad (2.40)$$

We have also introduced the normalised integrated Rabi frequency $\omega(\tau) = \frac{1}{W} \int_{-\infty}^{\tau} d\tau' |\Omega(\tau')|^2$, where the normalisation W is given by $W = \int_{-\infty}^{\infty} d\tau |\Omega(\tau)|^2$. As before, the time-dependent Rabi frequency $\Omega(\tau)$ describes the temporal profile of the control field and J_0 is a zero'th order Bessel function of the first kind.

The dimensionless coupling constant C is a measure for the strength of the memory interaction and uniquely determines the storage efficiency in the Raman limit $\Delta \gg \gamma$ where spontaneous emission can be neglected. Achieving sufficient coupling through a combination of large optical depth and high control pulse energy is therefore key to realising a strong memory interaction.

In the adiabatic limit we can find analytic expressions for the stored spin wave excitation B_{out} and the retrieved signal S_{out} , i.e.

$$\begin{aligned} B_{\text{out}}(z) &= \int_{-\infty}^{\infty} f(\tau) J_0 \left(2C \sqrt{(1 - \omega(\tau))z} \right) S_{\text{in}}(\tau) d\tau \\ S_{\text{out}}(\tau) &= f^*(\tau) \int_0^1 J_0 \left(2C \sqrt{\omega(\tau)(1 - z)} \right) B_{\text{out}}(z) dz, \end{aligned} \quad (2.41)$$

where $f(\tau) = C e^{iW\omega(\tau)/\Delta} \Omega(\tau) / \sqrt{W}$ is the normalised Stark-shifted Rabi frequency. Note that the normalised z -coordinate has been scaled by the ensemble length L so that the integral in the second of these equation is bounded by $z = 0$ and $z = 1$.

According to equation (2.35), the total memory efficiency can then be determined from

$$\eta_{\text{mem}} = \frac{N_{S_{\text{out}}}}{N_{S_{\text{in}}}} = \frac{\int_{-\infty}^{\infty} |S_{\text{out}}(\tau)|^2 d\tau}{\int_{-\infty}^{\infty} |S_{\text{in}}(\tau)|^2 d\tau} = \int_{-\infty}^{\infty} \left| \int_0^1 J_0 \left(2C \sqrt{\omega(\tau)(1 - z)} \right) f^*(\tau) B_{\text{out}}(z) dz \right|^2 d\tau. \quad (2.42)$$

2.2.3 Raman memory coupling

The energy of the control pulse can be approximated as $\mathcal{E}_c = I_c \mathcal{A} T_c$ where $I_c = \frac{1}{2} \epsilon_0 c \times |2E_c|^2$, \mathcal{A} is the cross-sectional area of the beam and T_c is the pulse duration. The peak control Rabi frequency Ω_{\max} can then be expressed in terms of \mathcal{E}_c :

$$\Omega_{\max} \approx \frac{\mathbf{d}_{23} \cdot \mathbf{v}_c}{\hbar} \times \sqrt{\frac{\mathcal{E}_c}{2\epsilon_0 c \mathcal{A} T_c}}. \quad (2.43)$$

If we approximate the shape of the control pulse as a top-hat function, the Raman coupling constant can be re-written as

$$C^2 \approx T_c d \gamma \times \left(\frac{\Omega_{\max}}{\Delta} \right)^2. \quad (2.44)$$

The optimal storage efficiency – predicted from the SVD of $K(z, \omega)$ in the Raman limit – rises sharply for small values of C , but plateaus for $C \sim 2$. This saturation effect is a result of stimulated scattering which becomes significant in the strong coupling regime^[170].

The ability to store multiple longitudinal modes in the Raman memory depends on the singular values λ_k of $K(z, \omega)$, i.e. the k^{th} mode can be stored with a maximum efficiency of $\eta_k = |\lambda_k|^2$. The multimode capacity N , however, only scales with the square root of the optical depth d so that for reasonable coupling strengths of $C \approx 1$, the Raman memory is essentially single-mode^[121]. The origin of this scaling can be identified as the scaling of the absorption bandwidth of an ensemble of 2-level atoms, each with a Lorentzian lineshape. It is in fact far more advantageous to store multiple input modes in spatially separate atomic ensembles, since in this case the multimode capacity scales linearly with d .

2.2.4 Beam splitter interaction

One of the distinguishing features of the Raman memory is the fact that its failure mode is transmission, i.e. any portion of the incident signal not mapped onto a spin wave coherence is transmitted through the ensemble, as the signal frequency is too far detuned from the atomic resonances to experience linear absorption. The memory can indeed be thought of as a beam splitter for optical and material modes. This analogy can be made rigorous in the Raman limit $\Delta \gg \gamma$ by defining a new set of variables σ and β and making the coordinate transformation $(\tau, z) \rightarrow (\omega, z)$:

$$\begin{aligned}\sigma(z, \omega)e^{-(W\omega+dz)/\Gamma} &= \sqrt{W} \frac{S(z, \tau)}{\Omega(\tau)} \\ \beta(z, \omega)e^{-(W\omega+dz)/\Gamma} &= B(z, \tau).\end{aligned}\tag{2.45}$$

This transformation is unitary provided that $\Gamma \approx -i\Delta$ is purely imaginary. The equations of motion (2.39) can now be simplified to the following symmetric set of coupled equations:

$$\begin{aligned}\partial_z \sigma &= -C\beta \\ \partial_\omega \beta &= C\sigma.\end{aligned}\tag{2.46}$$

The above equations are very similar to the quantum mechanical description of an optical beam splitter, except that the partial derivatives are with respect to both temporal and spatial coordinates. This is indicative of the presence of propagating optical modes and stationary material modes.

The solutions to these equations can be expressed in terms of Green's functions as

$$\begin{aligned}\sigma_{\text{out}}(\omega) &= \int_0^1 L(\omega, \omega') \sigma_{\text{in}}(\omega') d\omega' - \int_0^1 K(\omega, z) \beta_{\text{in}}(z) dz \\ \beta_{\text{out}}(z) &= \int_0^1 L(z, z') \beta_{\text{in}}(z') dz' + \int_0^1 K(z, \omega) \sigma_{\text{in}}(\omega) d\omega,\end{aligned}\tag{2.47}$$

where the Green's functions K and L relate the input signal to the stored excitation and the transmitted signal to the input field, respectively⁸.

From the coupled equations (2.46) it can be seen that $\partial_z(\sigma^\dagger \sigma) + \partial_\omega(\beta^\dagger \beta) = 0$. Integrating over z and ω gives the following conservation condition for the number of excitations involved in the memory interaction:

$$N_{\sigma, \text{out}} + N_{\beta, \text{out}} = N_{\sigma, \text{in}} + N_{\beta, \text{in}},\tag{2.48}$$

where $N_{\sigma, \text{in}}/N_{\sigma, \text{out}}$ are the numbers of incident/transmitted signal photons and $N_{\beta, \text{in}}/N_{\beta, \text{out}}$ are the numbers of initial/final spin wave excitations. This conservation law is based on the assumption of unitarity which is a good approximation in the Raman limit, provided that spontaneous emission can be neglected. Equation (2.48) further emphasises the analogy between the Raman memory and a beam splitter which is subject to a conservation condition of the same form⁹.

As a result of the unitarity of the above equations (and the persymmetry of the Green's functions), the SVDs of K and L share a common set of output modes and the squares of their singular values μ_j and λ_j add to unity. This is known as the *Bloch Messiah*

⁸Due to the symmetry between σ and β in the adiabatic equations of motion (2.46), we only require two distinct Green's functions K and L .

⁹Energy conservation between the input and output modes of a beam splitter ensures that photon numbers are preserved during the interaction.

Reduction^[220–222]. The Raman memory in the adiabatic limit is therefore formally equivalent to a light-matter beam splitter, described by

$$\begin{aligned} S_{\text{out},j} &= \mu_j S_{\text{in},j} + \lambda_j B_{\text{in},j} \\ B_{\text{out},j} &= \mu_j B_{\text{in},j} - \lambda_j S_{\text{in},j}, \end{aligned} \quad (2.49)$$

where $S_{\text{in/out},j}$ and $B_{\text{in/out},j}$ are the annihilation operators for a photon and spin wave excitation in the j^{th} input/output mode, respectively. The efficiencies, with which the j^{th} input mode is coupled to the j^{th} output mode is given by the quasi-reflectivity $\eta_j = \lambda_j^2$. Using the relations in (2.49), the Raman interaction for storage and retrieval can be illustrated pictorially as a system of two coupled beam splitters, as shown in Figure 2.1. Here, the spin wave output mode $B_{\text{out}}(t_1)$ from a storage event at t_1 becomes the input spin wave mode $B_{\text{in}}(t_2)$ for the retrieval interaction, provided that decoherence can be ignored on timescales of the storage time $(t_2 - t_1)$.

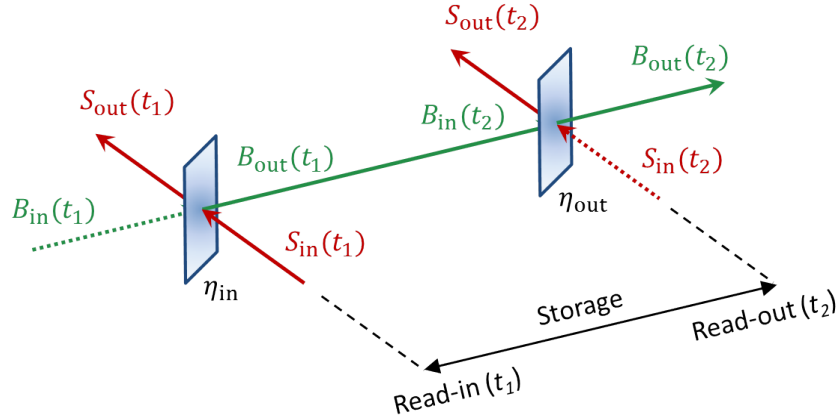


Figure 2.1: The Raman memory interaction as a beam splitter for optical and spin wave modes $S_{\text{in/out}}$ and $B_{\text{in/out}}$, respectively. Ideally, $B_{\text{in}}(t_1)$ during read-in is the vacuum contribution, likewise $S_{\text{in}}(t_2)$ during read-out. The storage and retrieval efficiencies are given by η_{in} and η_{out} .

The above discussion points towards an interesting application of Raman memories as systems which combine storage and beam splitter capabilities. By adjusting the storage

efficiency via the control pulse energy, one can furthermore realise a beam splitter with tunable reflectivity. This may be of particular interest in linear optical quantum computing architectures where beam splitters are one of the key components. An experimental demonstration of the Raman memory acting as a dynamically configurable network of beam splitters was reported in^[223], the results of which will be presented in Section 3.2.3.

2.3 Summary of Raman memory theory

By deriving an expression for the light-matter Hamiltonian, we have found a set of optical Bloch equations which describe the temporal evolution of the atomic populations and coherences in the rotating-wave approximation. Furthermore, we derived the propagation equation of the signal field using the paraxial and slowly-varying-amplitude approximations. Having defined continuous operators for the atomic polarisation and the collective spin wave coherence, we then arrived at a set of dynamical equations for the memory interaction.

The solutions to the memory equations of motion are given by a linear mapping, represented by a Green's function which relates input and output modes via a linear transformation. The shape of the optimal input mode and the associated memory efficiency can then be determined by performing an SVD of the Green's function. Moreover, the strength of the Raman interaction can be parametrised by the coupling constant $C \approx \sqrt{Wd}/\Delta$.

Finally, we have shown that the Raman memory – in the adiabatic limit – acts as a light-matter beam splitter coupling optical and material modes with a quasi-reflectivity given by the memory efficiency. Having established a theoretical framework for the Raman memory, we will next consider its experimental realisation in a warm ensemble of caesium atoms.

Chapter 3

The Raman memory experiment

“C.V. Raman was the first to recognize and demonstrate that the energy of a photon can undergo partial transformation within matter. I still recall vividly the deep impression that this discovery made on all of us...”

Albert Einstein

Since its first experimental demonstration in 2009^[174], the Raman memory has been successively refined, adapted and re-examined, ranging from incremental improvements in the experimental design to novel realisations in different systems. Most recently, the Raman memory has been successfully implemented in a 26- μm hollow-core photonic crystal fibre, paving the way for a fully integratable system^{[224] [225]}. This experiment was conducted by Michael Sprague as part of his PhD thesis, and is presently being led by PhD student Krzysztof Kaczmarek. Details of this work can be found in reference^[143].

3.1 Experimental setup

In the first part of this chapter, I will describe the core components of the experimental architecture, including the memory vapour cell, the preparation of the optical fields and the filtering and detection stages, as well as the basic experimental procedure. I will then give a brief overview of previously obtained results; in particular the initial demonstration of storage and retrieval of coherent states of light^[174], followed by experiments at the single-photon level^[226] and finally, the application of Raman memories as arrays of configurable beam splitters^[223].

3.1.1 Storage medium

The Raman interaction was implemented in the Λ -system of the caesium (^{133}Cs) D_2 line at 852 nm, shown in Figure 3.1 (a). The $F = 3$ and 4 states of the $6^2S_{1/2}$ ground-level hyperfine manifold serve as the initial and storage states, $|1\rangle$ and $|3\rangle$, respectively. These states, known as the caesium *clock states*¹, are dipole-decoupled^[228] and separated in frequency by 9.2 GHz due to the strong magnetic dipole coupling between the nuclear and electronic spins. The ground states are connected – via the D_2 line – to the upper $6^2P_{3/2}$ manifold which comprises the four hyperfine states $F = 2, 3, 4$ and 5. These states, each separated by approximately 200 MHz^[229], are not resolved in our memory experiments due to Doppler broadening at room temperature. We will therefore treat the excited state hyperfine manifold as the single excited state $|2\rangle$, as shown in Figure 3.1 (b).

Caesium was chosen as the storage medium, since it is the largest stable element in the group of alkali metals and thus has the longest diffusion time. Atomic diffusion determines the time it takes the atoms to escape the interaction region set by the optical beam diameter

¹The duration of the second is defined in terms of the hyperfine splitting of the clock states^[227].

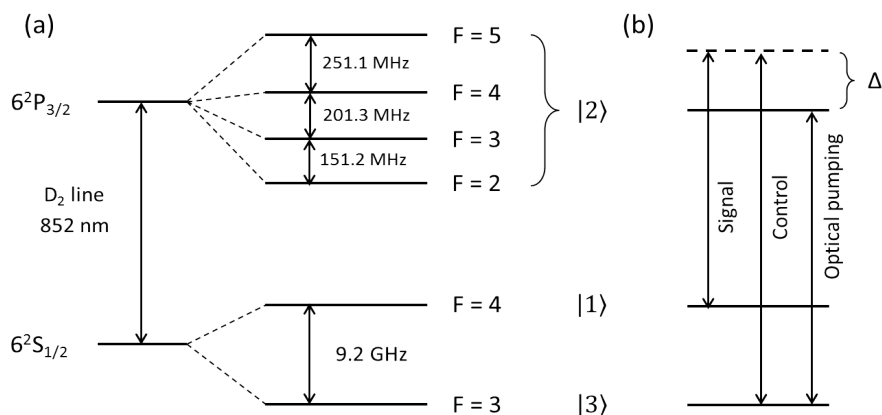


Figure 3.1: (a) Hyperfine structure of the caesium D_2 line (not to scale). The frequency splittings of the excited state are taken from^[230]. (b) Level diagram of the Raman memory. Signal and control fields are detuned by Δ from the excited state $|2\rangle$. The system is prepared in state $|1\rangle$ via optical pumping on the $|2\rangle \leftrightarrow |3\rangle$ transition.

and is one of the main decoherence mechanisms in vapour memories^{[174][175]}. In order to extend the diffusion time and thus the storage time, neon buffer gas was added to the cell. Collisions with neon atoms are spin-preserving and therefore do not affect the hyperfine Raman coherence, since no magnetic dipole interaction can occur between the caesium and neon nuclei². These collisions result in a random walk of the caesium atoms so that, for a typical beam diameter of $500 \mu\text{m}$, the diffusion time is on the order of several hundred microseconds^[170].

A further reason for the choice of caesium is the fact that its ground state hyperfine splitting of 9.2 GHz is the largest of the alkali metals. This splitting sets an upper limit on the possible memory bandwidth, since the signal and control pulses should not be so broadband as to address both the initial state and the storage state simultaneously.

The interaction strength of the Raman memory, as with other ensemble-based memories, is generally limited by the optical depth d , introduced in Section 2.2.1. By substituting the

²This is due to the fact that neon has zero spin and is therefore not associated with a magnetic moment.

definition of the coupling constant κ , the optical depth can be re-expressed in terms of the number density $n(T)$ of the atomic ensemble,

$$d = \frac{|\kappa|^2 L}{\gamma} = \frac{|\mathbf{d}_{12}^* \cdot \mathbf{v}_s|^2 \omega_s L}{2\gamma \epsilon_0 \hbar c} n(T). \quad (3.1)$$

The temperature-dependent number density can be found from the empirical vapour-pressure model given in reference^[229], using the ideal gas law $n(T) = P_v(T)V/RT$. The optical depth d is therefore directly proportional to the caesium vapour pressure in thermal equilibrium. In order to increase the vapour pressure and thus the strength of the light-matter interaction, the caesium cell was heated to a temperature of 70°C. While higher temperatures would further increase the optical depth, this would also reduce the optical pumping efficiency. This trade-off is the result of a phenomenon known as *radiation trapping*^[231] which describes the re-absorption of photons spontaneously emitted when atoms in the excited state relax to one of the ground states.

The caesium vapour, together with 20 Torr of neon buffer gas, was contained inside a cylindrical glass cell of length 7.5 cm and with anti-reflection-coated windows (Triad Technologies). The temperature of the cell was adjusted using an ohmic heating system of counter-wound wires which cancel magnetic fields arising from the heating current. Temperature regulation was implemented via a simple feedback mechanism using temperature sensors attached to various parts of the glass cell. In order to avoid condensation of caesium atoms on the cell windows, a *cold spot* was applied to a glass protrusion at the centre of the cell. This was achieved by directing room-temperature pressurised air onto the protrusion and ensuring that its temperature was a few degrees lower than at the windows.

The caesium cell was surrounded L with a single layer of μ -metal to provide magnetic

shielding and extend the storage time to the μs -regime. A de-Gaussing coil wrapped around the entire cell set-up was periodically used to minimise any residual magnetism.

3.1.2 Laser system

The control and signal pulses were both derived from a mode-locked Tsunami titanium-sapphire (Ti:Sa) laser which produced an 80-MHz train of pulses with a centre wavelength of 852 nm and a maximum output power of 1.5 W. The Ti:Sa laser was pumped with a 5-W frequency-doubled, continuous wave solid state laser (*Millenia Pro*) at 532 nm. An intracavity Gires-Tournois interferometer (GTI)³ was used to limit the pulse bandwidth to a FWHM of approximately 1.2 GHz, corresponding to a pulse duration of 360 ps. In addition, the laser frequency could be fine-tuned over a few tens of GHz by changing the free spectral range of the GTI. Coarse adjustment of the Ti:Sa frequency was achieved via a birefringent filter in the laser cavity.

3.1.3 Frequency lock

The frequency of the Ti:Sa was locked with respect to a frequency-stabilised diode laser which was also used for optically pumping the caesium atoms (see Section 3.1.5). Low-intensity pick-offs from the Ti:Sa oscillator as well as the diode laser were coupled into a scanning Fabry-Perot interferometer (SFPI). A frequency-offset lock, designed and implemented by Michael Sprague^[143], was used to actively stabilise the difference between the centre frequency of the Ti:Sa pulse and the diode frequency, as measured with the SFPI

³A GTI is a type of asymmetric Fabry-Perot etalon with one high-reflectivity surface, while the other surface is a partial reflector. It is used to provide negative group velocity dispersion (GVD) to compensate for the positive chirp of the oscillating pulse which results from the positive GVD of the intracavity optical components, in combination with self-phase modulation. By adjusting the spacing between the GTI plates, one can change the amount of negative GVD provided, as well as the range of frequencies over which the GVD is linear. The pulse bandwidth, for which lasing can occur, is therefore limited by the free spectral range (FSR) of the GTI.

(Thorlabs SA210-8B). This involved the application of a PID feedback signal to the piezo-electric driver which controls the GTI plate separation inside the Ti:Sa cavity. A data acquisition card (National Instruments PCI-6602) was used to generate this voltage signal and digitise the SFPI output. The PID input and output channels were controlled via a Labview program interfaced with MATLAB code which generated the error signal and feedback response. In order to lock the Ti:Sa frequency, the GTI voltage was initially swept over multiples of the SFPI free spectral range until both caesium resonances were identified by observing fluorescence emission from a reference cell. The Ti:Sa was then set to the same frequency as the diode laser which was locked to the $6^2S_{1/2}(F = 3) \leftrightarrow 6^2P_{3/2}$ transition (see Section 3.1.5). The feedback mechanism was engaged with a setpoint corresponding to zero offset, before changing the detuning to the desired value. This stabilised the Ti:Sa frequency to within approximately 100 MHz.

3.1.4 Signal and control field preparation

The signal and control fields were blue-detuned by 15 GHz from the excited state manifold. This detuning was chosen as it is sufficiently large to avoid linear absorption of the signal pulse, while also ensuring that the memory is operated in the adiabatic regime where $\Delta \gg \delta, \Omega$. Detuning the signal and control even further would, however, reduce the Raman coupling strength which is inversely proportional to the square of the detuning, as discussed in Section 2.2.2.

A Pockels cell (Quantum Technology QC) was used to set two pulse windows with which individual pulses for the storage and retrieval processes could be selected. A Pockels cell can be thought of as a voltage-controlled wave plate which can be used as an amplitude modulator when placed between two crossed polarisers. Only pulses whose polarisation has

been rotated by 90° are able to pass the second polariser. In order to achieve a full π -phase modulation, the Pockels cell was set up in a double-pass configuration. Rapid switching of the Pockels cell was possible with a rise time of < 1 ns. The timing and separation of the two Pockels cell windows were controlled by a digital delay generator (SRS DG535), synchronised to an 80-MHz timing trigger from the Ti:Sa oscillator. The repetition rate of the Pockels cell was set by dividing the 80-MHz pulse train by a variable factor, resulting in a repetition rate between 1 and 8 kHz. The extinction ratio between picked pulses and leakage was measured to be approximately 4000:1.

Pulse picking with the Pockels cell ensured that the signal and control fields were not present during the preparation and storage stages of the experiment. A single Pockels cell window would only allow storage and retrieval in consecutive time bins, thereby limiting the storage time to 12.5 ns. In order to increase the storage time beyond 12.5 ns, the size of the first window was reduced to select just a single write pulse, while the second window was used to delay the read pulses by up to several microseconds. The minimum window separation possible without degradation of the extinction ratio was approximately 100 ns.

As shown in Figure 3.2, the Ti:Sa beam was separated on a polarising beam splitter (PBS) into a strong, horizontally polarised control field arm and a weak signal field arm with vertical polarisation. The two-photon resonance condition of the Raman interaction requires that the frequencies of the control and signal beams are separated by the hyperfine splitting of the ground states $|1\rangle$ and $|3\rangle$. This was achieved with an electro-optic modulator (EOM) in the signal arm, generating sidebands shifted by 9.2 GHz from the fundamental laser frequency⁴. The red-detuned sideband corresponds to the signal frequency and was isolated from the other frequencies using an air-spaced planar Fabry-Perot etalon with a free

⁴Typical modulation efficiencies are on the order of a few percent.

spectral range (FSR) of 36.8 GHz. The FWHM of the etalon's transmission function was matched to the bandwidth of the laser pulse, resulting in a finesse of 24.5. By angle-tuning the etalon, the carrier frequency and the blue-detuned sideband could be attenuated by $\sim 10^{-2}$ relative to the signal frequency.

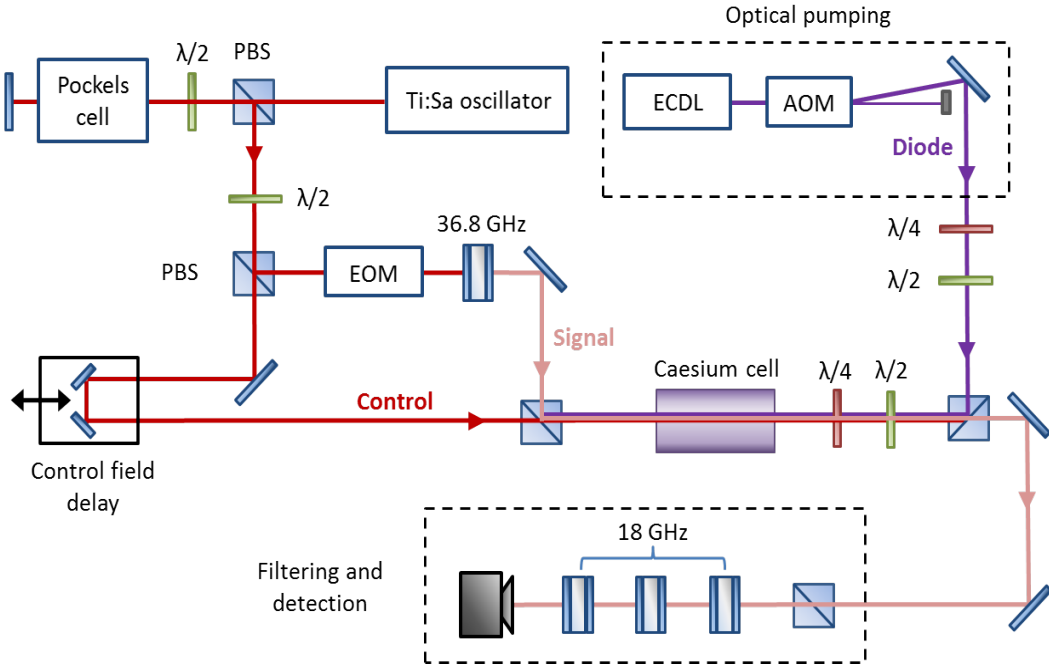


Figure 3.2: Experimental set-up of the Raman memory: Both the signal and control fields derive from a mode-locked Ti:Sa laser, producing 360-ps pulses at a repetition rate of 80 MHz. A Pockels cell was used to select pulses for storage and retrieval with a variable time delay. The signal field, red-detuned by 9.2 GHz from the control field, was generated using an electro-optic modulator (EOM) driven by a 9.2 GHz radio-frequency (RF) amplifier. Both the carrier frequency and the blue-detuned sideband were filtered using Fabry-Perot etalons (FSR 36.8 GHz). The signal and control pulses were overlapped on a polarising beam splitter (PBS) and focused into the caesium cell. The signal field was filtered with a polariser and three further etalons (FSR 18 GHz) and detected using a fast non-Geiger-mode APD. The optical pumping beam (ECDL) was introduced in a counter-propagating direction.

In order to ensure that only a single input signal was incident on the memory for each experimental cycle, the EOM was only switched on for the duration of the first read-in time bin. This was achieved with a fast RF switch which allowed switching times of approximately 3 - 4 ns. However, due to the finite response time of the EOM, a small residual signal

remained in the second time bin. When evaluating the memory performance for storage times of 12.5 ns, this residual was subtracted from the retrieved signal, so as not to overestimate the retrieval efficiency. The selection of signal and control pulses from the Ti:Sa pulse train is illustrated schematically in Figure 3.3.




Figure 3.3: A Pockels cell is used to select two pulse windows with variable delay from the 80-MHz train of Ti:Sa pulses. The signal pulse is prepared in the first time bin of Window 1 by switching the EOM. For storage and retrieval with non-consecutive control pulses, the size of the first window is reduced to pick a single write pulse.

The signal and control pulses were then recombined on a PBS and focused into the caesium cell. By introducing a suitable delay in the control field arm, the write pulse was temporally overlapped with the signal pulse to ensure optimal Raman coupling. This delay was set by maximising the interference visibility of pulses propagating along the signal and control beam paths using a polariser and a photodiode sufficiently slow to average the interference signal⁵.

The spatial overlap between signal and control modes was optimised by placing a flip mirror immediately before the memory cell. The diverted beams were then directed onto a beam splitter and imaged using two Thorlabs CMOS cameras positioned at either output port of the beam splitter. The distance between the flip mirror and the cameras was chosen

⁵The interference experiment was conducted by sending pulses of the same frequency along the signal and control arms. The response time of the photodiode was chosen to be large relative to the pulse duration.

such that the position of one camera corresponded to the front facet of the memory cell, while the other was equidistant from the centre of the interaction region. The pointing of the control field was then adjusted to overlap with the image of the signal beam on both cameras. This method ensured that the spatial overlap – a critical factor determining the memory efficiency – could be optimised in a repeatable manner⁶.

3.1.5 Optical pumping

Prior to any storage process, the atomic ensemble was prepared in the initial ground state $|1\rangle$ ($F=4$ of the $6^2S_{1/2}$ manifold) by optically pumping the atoms with an external cavity diode laser (ECDL) tuned into resonance with the $6^2S_{1/2}(F=3) \leftrightarrow 6^2P_{3/2}$ transition (see Figure 3.1 (b)). At room temperature, the two hyperfine ground states are equally populated, since their energy splitting $\mathcal{E}_{\text{HFS}} \ll k_B T$ where k_B is Boltzmann's constant and T is the temperature. After excitation from the $F=3$ ground state to the excited state, the atoms decay back to either ground state via spontaneous emission. With each cycle of excitation and decay, 50% of the population in the $F=3$ state is transferred to the $F=4$ state so that eventually, after multiples of the excited-state lifetime ($\tau_{D_2} \approx 30 \text{ ns}^{[229]}$), all atoms are spin-polarised.

The ECDL was directed into the caesium cell via the PBS used to separate the signal and control field polarisations immediately behind the vapour cell, as illustrated in Figure 3.2. By choosing a counter-propagating direction, the probability of detecting any residual ECDL noise was significantly reduced. For storage times exceeding 12.5 ns, it was necessary to switch off the diode laser during memory operation in order to avoid depletion of the

⁶When aligning the signal and control fields, the Ti:Sa was tuned to the signal frequency and the EOM was switched off so that the bright, unmodulated carrier component could be used instead of the weak sideband.

stored excitation. This was done by passing the ECDL through an acousto-optic modulator (AOM) so that its first diffraction order could be used as the pump beam. The AOM was then periodically switched off prior to each repetition of the memory cycle by synchronising the AOM-switch with the Pockels cell trigger using a digital delay generator.

The ECDL was set up in the well-known Littrow configuration^{[232] [233]}; the first diffraction order of an external grating (1800/mm) was coupled back into the laser diode, thereby providing frequency-selective optical feedback. The zeroth diffraction order was reflected out of the external cavity and used as the output beam. Precise frequency tuning could be achieved by adjusting the diode current as well as the piezo-electrically-controlled angle of the grating with respect to the cavity axis. The frequency of the diode laser was tuned to the $6^2S_{1/2}(F = 3) \leftrightarrow 6^2P_{3/2}$ transition by sweeping the diode current and grating angle, while observing resonant fluorescence in a caesium reference cell. The diode frequency was locked to this resonance by actively stabilising the driving current and grating angle via PID feedback^{[234] [235]}. The error signal was obtained from the Doppler-free absorption spectrum of a second caesium reference cell.

3.1.6 Filtering and detection

A crucial part of observing memory operation, especially at low signal photon numbers, is the extinction of the control field via polarisation, spatial and frequency filtering. The ability to do the former is a direct consequence of the fact that linearly polarised signal and control modes are only coupled by the Raman interaction if they are orthogonal to each other. In Section 4.1.1 we will examine this orthogonality condition more closely.

Polarisation filtering was achieved with a polarising beam splitter as well as a polarising beam displacer placed directly behind the memory cell. These provided an extinction ratio

of > 40 dB. The signal was then coupled into a single-mode fibre in order to ensure optimal spatial filtering. The remaining control field was further attenuated with a series of air-spaced Fabry-Perot etalons, angle-tuned for maximum transmission of the signal frequency. Three etalons, each with an FSR of 18.4 GHz and an extinction ratio of 10^{-2} , were used to frequency-filter bright coherent input states of several thousand photons. For input signals at the single-photon level, it was necessary to add a further etalon (FSR = 103 GHz) in a double-pass configuration in order to achieve a sufficient degree of filtering (see Section 5.2.1).

In the case of bright coherent input states, the signal was detected with a high-sensitivity avalanche photodiode (Menlo Systems APD210) connected to a GHz-bandwidth oscilloscope (LeCroy WavePro 7100). For single-photon level measurements, the filtered signal – coupled into a multimode fibre – was detected using a single-photon counting module (PerkinElmer SPCM-AQRH-R-FC) operating in Geiger mode. Detection events were counted with a *Xilinx* field-programmable gate array (FPGA), developed by PhD student Justin Spring^[236]. In addition, arrival time histograms were recorded with a time-to-amplitude converter (TAC) and a multi-channel analyser (MCA).

The average number of signal photons incident on the memory could be estimated by correcting the number of detection events N_{det} for all transmission losses T_{tot} between the front of the vapour cell and the detector, and for the APD detection efficiency $\eta_{\text{det}} = 0.45 \pm 0.05$. Hence, the estimated input photon number N_{in} is given by

$$N_{\text{in}} = \frac{N_{\text{det}}}{T_{\text{tot}} \cdot \eta_{\text{det}}}. \quad (3.2)$$

Here, we have defined the total transmission $T_{\text{tot}} = T_{\text{SMF}} \cdot T_{\text{ETL}} \cdot T_{\text{MMF}}$, where $T_{\text{SMF}} \approx 70\%$ is the coupling efficiency into the single-mode fibre, $T_{\text{ETL}} \approx 10\%$ is the total transmission

through the series of etalons and $T_{\text{MMF}} \approx 90\%$ is the coupling efficiency into the final multimode fibre.

3.2 Previous results

The first experimental realisation of the Raman memory was demonstrated in 2009 by PhD student Klaus Reim. In the following sections, I will briefly outline some of the key results of this pioneering investigation.

3.2.1 Storage of bright coherent states

The Raman memory was initially tested using bright coherent signal pulses of approximately 10^5 photons for a storage time of 12.5 ns and with a total efficiency of up to 15%^[174]. On this timescale, the spin wave excitation is unaffected by decoherence mechanisms so that the memory efficiency can be taken as a direct indicator of the intrinsic Raman coupling strength. The dependence of the total memory efficiency on the control pulse energy was measured and compared to theoretical predictions based on equation (2.42), as shown in Figure 3.4.

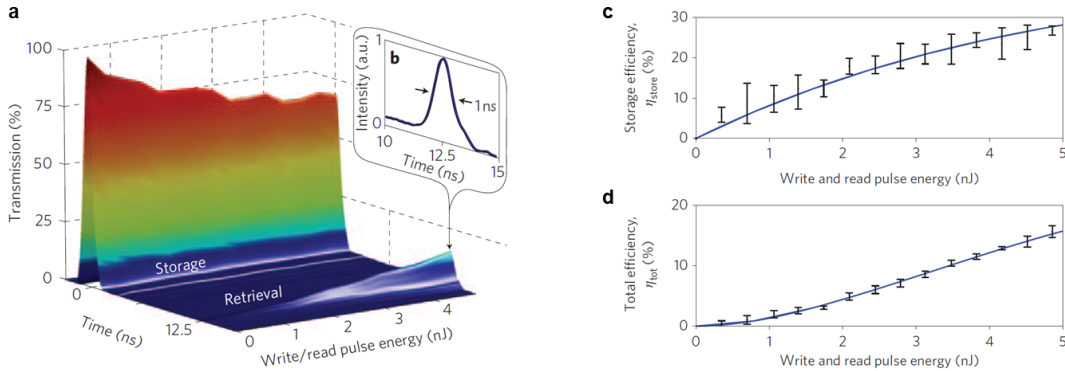


Figure 3.4: **a** Experimental data for storage and retrieval after 12.5 ns as a function of control pulse energy^[174]. **b** Temporal profile of the retrieved signal with a FWHM of 1 ns, limited by the detector response time. **c** Storage efficiency. **d** Total memory efficiency. Experimental data is shown in black with error bars indicating the standard deviation from a set of 100 sample measurements. Solid lines represent the theoretically predicted dependence of the efficiency on the control field energy.

The storage and retrieval efficiencies, η_{stor} and η_{retr} , were limited by the available control pulse energy of ~ 10 nJ. In principle, forward retrieval places an upper bound on η_{retr} due to re-absorption of the retrieved signal field, as it propagates through the atomic ensemble. Optimal read-out can indeed only be achieved through backward retrieval which maximises the spatial overlap between the spin wave generated during storage and the spin wave mode for which the retrieval process is optimised⁷. However, unless the frequencies of the signal and control fields are identical, the spin wave acquires a spatially varying phase which imposes a non-collinear phase-matching condition. In order to ensure momentum conservation during read-out, phase-matched backward retrieval therefore requires a small angle between the signal and control beam directions^[237].

The phase coherence of the memory interaction was tested by interfering the retrieved pulse with a delayed copy of the original signal in a Mach-Zender-type interferometer, yielding a corrected fringe visibility⁸ of $82.7 \pm 0.9\%$, as shown in Figure 3.5.

⁷In the case of forward retrieval, the spin wave mode corresponding to optimised storage is poorly overlapped with the optimal spin wave mode for retrieval^[237].

⁸The visibility has been normalised by the visibility of the interferometer to correct for imperfect alignment

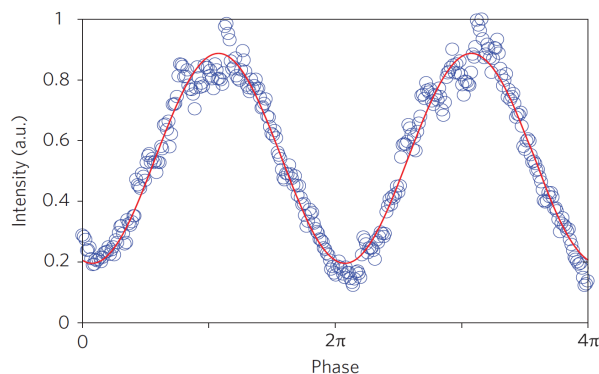


Figure 3.5: Interference of stored and retrieved signals, measured by scanning the delay of a copy of the input signal^[174]. The data (blue circles) is fitted with a least-square fit (red line), giving a visibility of $82.7 \pm 0.9\%$.

This value is consistent with the maximum visibility predicted from a theoretical model which takes into account the expected pulse distortion of the retrieved signal field due to the A.C. Stark shift arising from the strong control field^[226].

3.2.2 Storage at the single-photon level

Following these initial results, storage and retrieval of weak coherent light pulses at the single-photon level was demonstrated for storage times of up to $4 \mu\text{s}$ and with a total efficiency of 30% ^[226]. Figure 3.6 shows the memory data for an input signal of 1.6 photons per pulse. The retrieved signal field is followed by an exponentially decreasing tail which can be attributed to fluorescence noise from the excited state manifold. This is the result of the collision-induced excitation of atoms from the ground state to the excited state in the presence of the control field. The decay time of $\tau \sim 32 \pm 2 \text{ ns}$, determined from a least-squares fit of the exponential tail (red curve), is in good agreement with the excited state lifetime of the caesium $D_2 \ 6^2S_{1/2} \rightarrow 6^2P_{3/2}$ transition (30.4 ns)^{[229] [238]}. As this timescale is two orders of magnitude longer than the duration of the retrieval process, fluorescence and phase fluctuations.

noise could be easily removed by time-gating the detection.

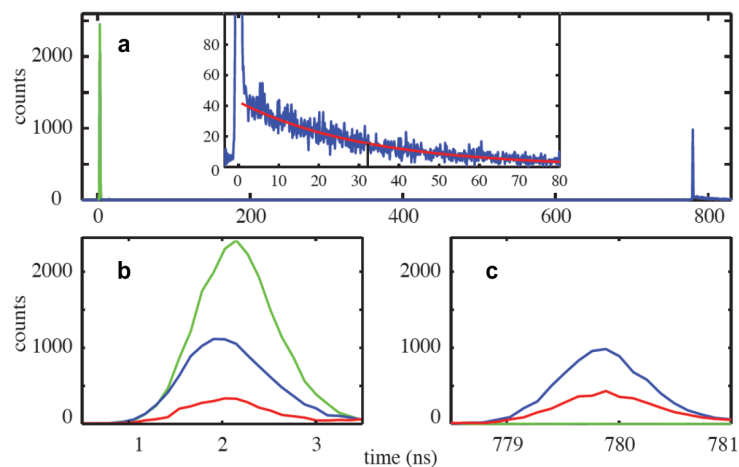


Figure 3.6: **a** Storage and retrieval after 780 ns^[226]. Inset: Detail of the retrieved signal, showing an exponentially decreasing tail due to fluorescence noise. **b** Read-in time bin. **c** Read-out time bin. TAC histograms, accumulated over 360000 runs, show the incident signal in the absence of the control field (green), the transmitted/retrieved signal (blue) and noise (red).

In addition to fluorescence noise, an instantaneous noise contribution was measured in the absence of the signal field. This unconditional noise floor of 0.25 photons per pulse can be ascribed to a four-wave mixing process seeded by spontaneous anti-Stokes scattering^[239]^[240]. A detailed discussion of this noise source will be presented in Section 5.4.

The memory lifetime was extended to approximately 1.5 μ s by adding μ -metal magnetic shielding around the vapour cell, thus reducing the effects of spin wave dephasing caused by residual magnetic fields. Nevertheless, it is understood that magnetic dephasing remains the dominant limitation on the memory lifetime, since decoherence as a result of atomic diffusion is expected to occur on timescales of hundreds of microseconds^[175]. The measurements of memory efficiency as a function of storage time (Figure 3.7) can be compared to the predictions of a simple model which I will describe below.




Figure 3.7: Memory efficiency as a function of storage time^[226]. Experimental data is shown in blue, with error bars representing the standard deviation of the mean. The solid line is the theoretical prediction based on magnetic dephasing of the spin wave (see below).

Magnetic dephasing

Magnetic dephasing can be modelled by considering the phase accumulated by each Zeeman state $|F_j, m_j\rangle$ due to Larmor precession of the atomic spins under the influence of a static magnetic field. To see this, we introduce the spin wave operator $\Xi = \alpha I + \beta \Sigma$ where I is the identity operator. The transition operator Σ is given by

$$\Sigma = \sum_{m_i=-F_i}^{F_i} \sum_{m_f=-F_f}^{F_f} C(m_i, m_f) |F_i, m_i\rangle \langle F_f, m_f|. \quad (3.3)$$

where $C(m_i, m_f)$ is the coupling coefficient between the initial and final Zeeman states, $|F_i, m_i\rangle$ and $|F_f, m_f\rangle$, respectively.

Larmor precession in the presence of a magnetic field B can be described by the operator

$U = R^\dagger ER$, where

$$\begin{aligned}
 R = & \sum_{m_i=-F_i}^{F_i} e^{i(Y_i \sin \phi - X_i \cos \phi)\theta} |F_i, m_i\rangle \langle F_i, m_i| \\
 & + \sum_{m_f=-F_f}^{F_f} e^{i(Y_f \sin \phi - X_f \cos \phi)\theta} |F_f, m_f\rangle \langle F_f, m_f|
 \end{aligned} \tag{3.4}$$

rotates the quantisation axis onto the direction of the magnetic field, parametrised by the azimuthal angles (θ, ϕ) and

$$\begin{aligned}
 E = & \sum_{m_i=-F_i}^{F_i} e^{im_i g_i \mu_B B t} |F_i, m_i\rangle \langle F_i, m_i| \\
 & + \sum_{m_f=-F_f}^{F_f} e^{im_f g_f \mu_B B t} |F_f, m_f\rangle \langle F_f, m_f|
 \end{aligned} \tag{3.5}$$

accounts for the phase accumulation of each Zeeman level. Here, $X_{i,f}$ and $Y_{i,f}$ are the x and y components of the spin angular momentum operators and $g_{i,f}$ are the Landé g -factors for the states $F_{i,f}$.

The retrieval efficiency as a function of storage time $\eta(t)$ is proportional to the square of the expectation value of Σ and can be expressed as a sum of contributions from each initial Zeeman state,

$$\eta \propto \sum_{m_i=-F_i}^{F_i} p_{m_i} |\langle F_i, m_i | U^\dagger \Sigma U | F_i, m_i \rangle|^2 \tag{3.6}$$

where $p_{m_i} = 1/(2F_i + 1)$. The efficiency $\eta(t)$, multiplied by a scaling factor, is plotted in Figure 3.7 together with the experimental data from which the magnetic field $B = 0.13 \pm 0.05$ Gauss can be determined using a least-squares fit. This value is consistent with a small residual component arising from the earth's magnetic field.

In principle, the spin coherence should undergo periodic rephasings in the presence of a

homogeneous magnetic field due to the discrete number of Zeeman states. However, since the earth's magnetic field is very weak, the rephasing time $T_{\text{rephase}} \sim 1/|\mathbf{B}_{\text{earth}}|$ is significantly longer than the timescales of inhomogeneous decoherence mechanisms, including atomic collisions and diffusion. By applying a sufficiently strong external magnetic field, T_{rephase} can be reduced to $\sim 10\mu\text{s}$, thus allowing magnetic rephasing of the spin wave to be observed. This has been demonstrated experimentally by Klaus Reim^[241]. While further investigation of the rephasing effect will be necessary, it may provide an interesting route towards extending the memory lifetime.

3.2.3 Multi-pulse addressing

As outlined in Section 2.2.4, the Raman interaction can be reduced to a sum of independent beam splitter transformations via the Bloch-Messiah reduction^[172], provided that the detuning of the optical fields is sufficiently large to validate the Raman approximation. In this limit, spontaneous emission can be neglected and the light-matter interaction becomes unitary^[170]. The lack of resonant absorption ensures that unstored signal photons are transmitted through the memory medium without attenuation. Similarly, the stored spin wave excitation is preserved on timescales short compared to the decoherence time, since no population is transferred to the excited state. This beam-splitter functionality was tested experimentally by converting the Raman memory into an array of dynamically configurable beam splitters^[223].

Following initial signal storage, the memory was addressed with a train of multiple control pulses, selected with the second window of the Pockels cell. The remaining signal could then be read out over multiple control pulses, yielding an overall retrieval efficiency arbitrarily close to unity.




Figure 3.8: Multi-pulse retrieval^[223]. The control pulses for (a) storage (b)-(e) retrieval over up to four consecutive time bins are shown in blue (middle panel). Red lines correspond to transmitted/retrieved signal pulses and thick grey lines show the theoretical prediction (Note that the signal retrieved after quadruple read-out is too weak to be distinguished from noise.). The top panel shows the corresponding beam splitter networks for successive memory interactions coupling optical and spin wave modes.

The results are plotted in Figure 3.8, showing multiple retrieval of the spin wave with up to four consecutive control pulses, alongside the analogous beam splitter networks. The combined retrieval efficiency of these four read-out events was approximately 95%. This value agrees with the theoretically predicted efficiencies in each read-out bin, as determined from the model developed in Section 2.2.2. It should be noted that in the case of multi-pulse retrieval, the shape of the control read pulse $\Omega(\tau)$ is given by a train of pulses. The temporal profiles S_{in} and $\Omega(\tau)$ were determined experimentally in order to solve the retrieval mapping $B_{\text{in}} \rightarrow S_{\text{out}}$, given by equation (2.32)⁹.

As discussed in Section 2.2.4, a set of optical and spin wave modes, independently coupled by effective beam splitter interactions, can be found by performing an SVD of the Green's function K . The reflectivity of each of these interactions is thereby determined by the control field energy. This was tested by adjusting the Pockels cell timing in order to

⁹For a storage time of $T = 900$ ns, the decoherence of the spin wave was accounted for by setting $\eta_T = 0.7$.

vary the energy in each control pulse of the read sequence. In this manner, the distribution of the retrieved signal over multiple temporal modes could be dynamically adjusted, as shown in Figure 3.9. The ability to manipulate partial storage and retrieval highlights the configurability of the Raman memory beam splitter network.

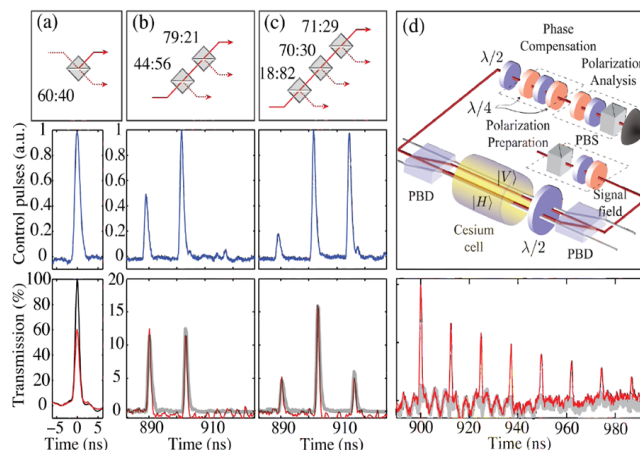


Figure 3.9: Shaped retrieval (colour coding as in Figure 3.8). (a) Storage event. (b) Balanced retrieval achieved by adjusting the pulse energies of the control field. (c) Pyramidal retrieval. (d) Phase coherence measurement. Bottom panel: Constructive (red) and destructive (grey) interference between two pulse trains retrieved from independent memories. Top panel: Experimental layout. See Section 4.2 for a detailed description.

Finally, the phase coherence of the Raman interaction across multiple time bins was confirmed by interfering the retrieved signals from two independent memories contained within a polarisation interferometer. The results are presented in Figure 3.9 (d), showing almost perfect destructive interference in all read-out time bins.

The polarisation interferometer depicted in the top panel of Figure 3.9 (d) will be discussed in more detail in the next chapter, where we will examine the ability of the Raman memory to preserve the phase coherence of arbitrary polarisation states of light.

Chapter 4

Polarisation storage

Photonic qubits can be encoded in several different degrees of freedom of the electromagnetic field, including photon number^[242–244], spatial mode^[245], time bin^{[246][247]}, frequency^[248] and polarisation^[249]. The polarisation degree of freedom in particular allows qubits to be transmitted in a single spatial and temporal mode. Experimentally, polarisation-encoded photonic qubits provide a convenient choice, since polarisation states can be easily prepared, manipulated and analysed^[250].

The faithful storage and retrieval of polarisation-encoded information is a key requirement for realising a scalable photonic network with polarisation qubits. In the following chapter, we will examine the Raman memory’s ability to preserve the polarisation state of a stored signal using weak coherent states as inputs.

The experimental design of the polarisation memory is based on initial work by Klaus Reim, as described in his PhD thesis^[241]. My contribution consisted in performing the subsequent experiments presented below, alongside PhD students Duncan England and Patrick Michelberger.

4.1 Storage of polarisation states

As discussed in Section 2.2.3, the Raman memory is essentially a single-mode memory^[121] which can only efficiently store one polarisation mode per spin wave¹. It is therefore not possible to store polarisation-encoded information in a single Raman memory.

The Raman interaction is in fact polarisation-sensitive, i.e. signal and control fields are only coupled by the interaction if they have orthogonal linear polarisations or the same circular polarisation^[170]. This condition will be derived in the following section.

4.1.1 Polarisation selection rules

In order to understand the polarisation selectivity of the Raman interaction, we need to consider the selection rules determining the allowed transitions between the initial atomic state $|F_i, m_i\rangle$ and the final state $|F_f, m_f\rangle$, via an intermediate state $|F_m, m_m\rangle$ from the $6P_{3/2}$ manifold (see Figure 3.1). As each photon only carries one quantum of angular momentum, only the two central hyperfine states $F = 3$ and $F = 4$ can act as intermediate states.

In the absence of a magnetic field, the magnetic quantum numbers m of the $(2F + 1)$ Zeeman sublevels can be defined in terms of angular momentum projections along an arbitrary quantisation axis, here taken to be the z -axis. A field polarised linearly along this quantisation axis, known as π -polarised light, cannot induce a change in magnetic quantum number ($\Delta m = 0$). This can be verified by integrating the angular part of the dipole matrix element for an electric field directed along the z -axis and noting the cylindrical symmetry of this integral. Indeed, the emergence of polarisation selection rules can be understood as a consequence of symmetry considerations^[251].

¹In principle, the number of modes which can be effectively stored in the memory is given by the number of non-negligible singular values in the SVD of $K(z, \tau)$.

For a π -polarised control field the possible intermediate states of the Raman interaction are therefore restricted to $|F_m = 3, m_m\rangle$ and $|F_m = 4, m_m\rangle$. In this case the emitted field cannot be π -polarised as well, since the transition pathways via the two intermediate states with $F_m = 3$ and $F_m = 4$ interfere destructively². This can be shown by evaluating the corresponding scattering amplitude $A_{\pi\pi}$:

$$A_{\pi\pi}(F_i, m_i) = \sum_{F_m=3,4} d(J = 1/2, F_i, m_i \rightarrow J = 3/2, F_m, m_m = m_i) \quad (4.1)$$

$$\times d(J = 3/2, F_m, m_m = m_i \rightarrow J = 1/2, F_f \neq F_i, m_f = m_i) = 0$$

where the dipole matrix elements d can be determined from the Clebsch-Gordan coefficients. Hence, a vertically polarised signal field can only be stored using a horizontally polarised control field, or vice versa.

A similar argument shows that for circularly polarised fields propagating along the z -axis, the scattering amplitude is only non-zero if both the signal and control fields are either σ_+ - or σ_- -polarised. Here, the selection rule for the magnetic quantum number is $\Delta m = \pm 1$, again for reasons of symmetry^[251]. As we will see in Section 5.6.2, this unfortunately precludes the ability to suppress noise emission during memory operation using polarisation selection rules.

4.2 The polarisation interferometer

Because of the polarisation selectivity as well as the single-mode nature of the Raman memory, polarisation storage could only be achieved by decomposing the input state into its orthogonally polarised components. Each of these components was then mapped onto a

²This result assumes that the detuning Δ of the optical fields is large compared to the splitting of the intermediate excited states.

separate atomic ensemble. The resulting two independent memories were contained within the same caesium cell and addressed with collinear, orthogonally polarised signal and control fields in a passively stable interferometric set-up^{[252][249]}, known as a Jamin-Lebedev-type interferometer. The vapour cell was placed in between two polarising beam displacers (PBD) which deflect only the horizontally polarised component of the incident light, with a polarisation extinction of approximately 10^{-5} . As a result, orthogonally polarised components emerge as two parallel output beams, spatially separated by 8 mm. The geometry of this dual-rail polarisation memory is shown in Fig. 4.1.

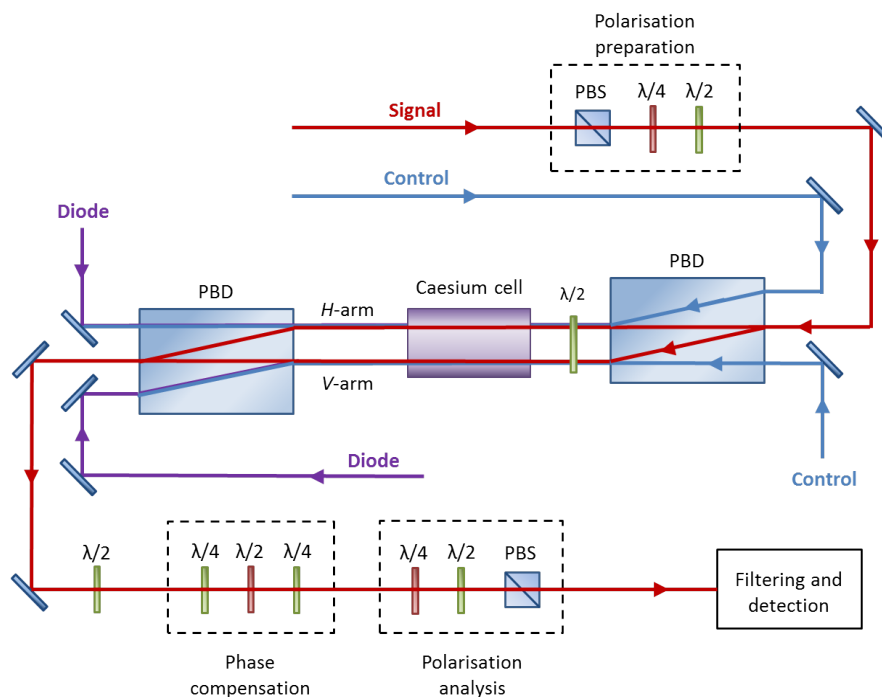


Figure 4.1: Polarisation interferometer: Storage of an arbitrary input state is achieved by mapping the horizontally and vertically polarised components of the signal field onto separate ensembles within the same vapour cell.

The input signal was prepared in an arbitrary polarisation state $\cos \theta |H\rangle + e^{i\phi} \sin \theta |V\rangle$ using a PBS, followed by a $\lambda/2$ and a $\lambda/4$ waveplate in front of the PBD set-up. Both the control field and pump beam were split into horizontally and vertically polarised compon-

ents which were directed into the memory such that each component overlapped with an orthogonally polarised signal field arm.

This configuration establishes two parallel memory arms for the independent storage of horizontally and vertically polarised signals. A $\lambda/2$ waveplate after the first PBD was set to rotate each polarisation by 90° such that the signal beams from both memory arms would recombine when exiting the second PBD, while the control field paths would be spatially separated. Subsequent polarisation analysis of the combined signal was performed using calibrated $\lambda/2$ and $\lambda/4$ waveplates as well as a Glan-Laser polariser aligned parallel with the preparation PBS.

4.3 Quantum process tomography

In order to characterise the dynamics of the memory interaction fully, we can use the polarisation interferometer described above to perform a procedure known as *quantum process tomography* (QPT) on the system^[253–255]. A quantum process is a completely positive map relating the density matrices ρ_{in} and ρ_{out} of the initial and final states, i.e.

$$\rho_{\text{out}} = \sum_{m,n} \chi_{mn} \Gamma_m \rho_{\text{in}} \Gamma_n \quad (4.2)$$

where χ_{mn} is the process matrix which is positive Hermitian^[3]. The set of operators Γ_m form a complete orthonormal basis for the space of density matrices and for single qubit tomography it is convenient to express Γ_m in the Pauli basis, i.e. $\Gamma_{0,1,2,3} = \{I/\sqrt{2}, X/\sqrt{2}, Y/\sqrt{2}, Z/\sqrt{2}\}$. QPT is used to reconstruct the process matrix χ_{mn} which contains complete information about the dynamics of the system.

The density matrix ρ of a single qubit can be expanded in terms of the orthonormal

basis set as

$$\rho = \frac{\text{tr}(\rho)I + \text{tr}(X\rho)X + \text{tr}(Y\rho)Y + \text{tr}(Z\rho)Z}{2}, \quad (4.3)$$

where quantities of the form $\text{tr}(X\rho)$ are equal to the expectation value of the observable X . By measuring X a large number of times, m , with outcomes $x_1, x_2, x_3, \dots, x_m$, we can obtain an estimate of the expectation value, given by the average $\sum_i x_i/m$. Similarly, estimates of Y and Z can be determined through repeated measurement. If the sample size is sufficiently large, we can use these averages to obtain a good estimate of ρ via a process known as *quantum state tomography* (QST)^[3].

We are now in a position to characterise the quantum process (4.2) by choosing a set of input states and performing QST on each of the output states. Provided the input states form a complete basis set spanning the input Hilbert space^[253], these measurements contain all accessible information about the quantum process. Hence, QPT allows the prediction of ρ_{out} for an arbitrary input state.

In order to quantify the effect of the memory interaction on the stored state, the process matrix χ is determined for two experimental conditions; χ_{on} corresponds to the memory being ‘switched on’ in the presence of the control field, while χ_{off} is obtained when the control field is blocked. The process fidelity $F = (\text{tr}\sqrt{\sqrt{\chi_{\text{on}}}\chi_{\text{off}}\sqrt{\chi_{\text{on}}}})^2$ can then be computed to characterise the memory performance in terms of the overlap between the stored and retrieved quantum states^[102]. A fidelity of $F = 1$ corresponds to an ideal memory process which exactly returns the original state and is equivalent to the identity operation. In contrast, $F < 1$ indicates that the input state has been altered during the memory interaction. The process fidelity can therefore be understood as a measure of the memory’s ability to preserve the coherence of the stored polarisation state.

As discussed in the previous chapter, potential sources of decoherence include magnetic dephasing, collisional spin-flips and atomic diffusion out of the interaction region.

4.4 Experimental procedure

Process tomography (PT)³ can be performed by preparing the signal field in a set of polarisation states which span the Poincaré sphere and form a mutually unbiased bases (MUB)^[256–258]. Since any arbitrary polarisation state can be constructed from these MUB states, it is sufficient⁴ to demonstrate storage and retrieval of these states in order to verify full polarisation storage^[259]. The output state for each of the prepared input polarisations is then analysed by projective measurements along the chosen set of states. This is equivalent to performing state tomography for each input polarisation. From these measurements, the process matrix can then be reconstructed through maximum likelihood estimation^{[260][261]}.

In practice, process tomography is commonly performed with an overcomplete set of six polarisation states, as this redundancy tends to add robustness to the estimation method. Hence, the input signal was sequentially prepared in $|H\rangle$, $|V\rangle$, $|D\rangle = \frac{1}{\sqrt{2}}(|H\rangle + |V\rangle)$, $|A\rangle = \frac{1}{\sqrt{2}}(|H\rangle - |V\rangle)$, $|R\rangle = \frac{1}{\sqrt{2}}(|H\rangle - i|V\rangle)$, $|L\rangle = \frac{1}{\sqrt{2}}(|H\rangle + i|V\rangle)$, where $|H\rangle$ and $|V\rangle$ are the horizontal and vertical polarisation states, respectively. The resulting set of 36 measurements was repeated 500 times from which error bars were determined using a Monte Carlo simulation⁵.

For each of the six input polarisations, data sets were recorded on a fast oscilloscope (*LeCroy* LCwave 7300) with the following experimental settings:

³Since the input signal is a coherent state, rather than a single photon, we will refrain from referring to it as *quantum* process tomography.

⁴This is a valid argument due to the linearity of the memory process.

⁵The algorithm used to perform the reconstruction of the process matrix was supplied by Nathan Langford^[253].

1. *sd*: signal and diode present/control field blocked, corresponding to a measurement of the input signal

2. *scd*: all fields present, allowing full memory operation

3. *cd*: signal blocked, corresponding to a noise measurement

A *LabView* script was used to automate data acquisition by controlling a pair of shutters with which the signal and control fields could be independently blocked between measurement runs. The total memory efficiency in read-out time bin i , η_{tot}^i , was determined according to

$$\eta_{\text{tot}}^i = \frac{A_{scd}^i - A_{cd}^i}{A_{sd}^0 - A_{cd}^0} \quad (4.4)$$

where A_j^0 and A_j^i are numerically integrated pulse traces in the read-in time-bin and the i^{th} read-out time-bin, respectively. The subscript $j \in \{sd, scd, cd\}$ denotes the experimental setting.

Throughout these measurements, the memory efficiency in each arm of the polarisation interferometer had to be carefully balanced by adjusting the control field splitting. When storing polarisation states formed from a superposition of horizontally and vertically polarised components, any imbalance between the memory arms would otherwise have manifested itself as an artificial rotation of the retrieved state.

Previous PT measurements showed that the process purity⁶, defined as $P = \text{tr} \{\chi^2\}$, was severely degraded as a result of phase fluctuations in the interferometer. These increased with the vapour cell temperature, leading to purities of only 64% and 46% at 62.5°C and 65°C, respectively^[226]. This phase instability was caused by air turbulence due to the

⁶A perfectly coherent process subject to unitary dynamics yields $P = 1$, while a completely incoherent process corresponds to a purity of 1/4.

temperature gradient between the heated cell and the room-temperature environment. By surrounding the beam paths inside the interferometer with plastic tubing, it was possible to reduce the effects of such external perturbations. In addition, phase differences between the interferometer arms were corrected using a set of three waveplates before the polarisation analysis stage, i.e. $\lambda/4$, $\lambda/2$ and $\lambda/4$ waveplates provided Berry phase compensation^{[253][262]}. The fast axis of each $\lambda/4$ waveplate was set to $\pm 45^\circ$ to the horizontal so that the linear polarisation of the incident signal is unaffected. However, the signal's polarisation state is taken around a closed loop on the Poincaré sphere, thereby acquiring a geometric phase shift depending on the solid angle subtended by the loop. This phase shift can be adjusted using the $\lambda/2$ waveplate in order to compensate for phase differences between the interferometer arms. Berry-phase compensation therefore ensures that input signals consisting of superpositions of $|H\rangle$ and $|V\rangle$ are faithfully reproduced. An additional $\lambda/2$ waveplate compensated for the 90° polarisation rotation inside the interferometer, as shown in Figure 4.1. While the polarisation interferometer was designed to be passively stable, it was necessary to check the settings of the phase compensation set-up regularly during the course of the measurements. This was done by preparing a diagonally polarised input signal $|D\rangle$ and minimising the transmitted field along $|A\rangle$, while blocking the control field.

The following measurements were made using coherent input states with average photon numbers of 1000 - 10000 per pulse. While the results are equally applicable for memory operation at the single-photon level^{[216][263]}, a few modifications would have to be made to the interferometric set-up. For instance, small waveplates would have to be inserted in each arm of the interferometer to compensate for the birefringence of the cell walls. With the present set-up, this birefringence gives rise to a small rotation of the control field polarisation which would cause leakage of control photons through the polarisation filtering. Currently,

these additional noise counts would be prohibitive for measurements at the single-photon level. In addition, active stabilisation of the interferometer would become necessary due to the long measurement times needed to build up single-photon statistics.

4.5 Results of the process tomography

Following the procedure outlined above, full process tomography was performed for a range of different storage times. The process matrix χ_{off} , determined from the input signal measured in the absence of the control field (sd), allows us to characterise the bare interferometer. As shown in Figure 4.2 (a), χ_{off} consists almost entirely of the identity I with an associated process purity of $P_{\text{off}} = \text{tr} \{ \chi_{\text{off}}^2 \} = 0.989$, confirming that the polarisation interferometer is phase-stable over the course of the measurements. Figure 4.2 (b) shows a typical example of the process matrix χ_{on} for states retrieved after a storage time of 750 ns, giving rise to a corresponding purity of $P_{\text{on}} = \text{tr} \{ \chi_{\text{on}}^2 \} = 0.846$. The dominance of the identity transformation demonstrates that the memory is indeed highly capable of preserving polarisation coherence.

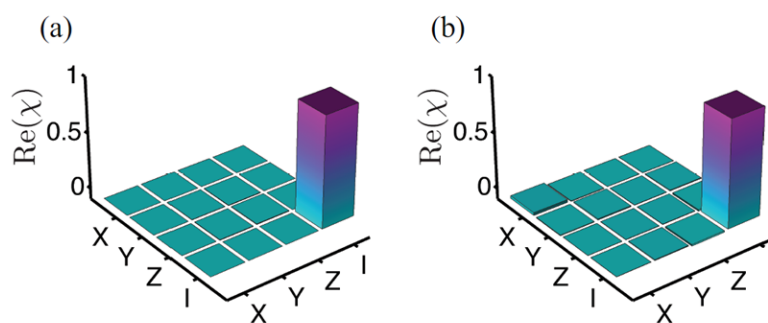


Figure 4.2: Reconstructed process matrices measured in the absence of the control field (a) and with the control field present, for pulses retrieved after a storage time of 750 ns (b).

The process fidelities, alongside the corresponding memory efficiencies, are plotted in

Figure 4.3 for a number of different storage times. The data sets were Fourier-filtered by applying a top-hat filter function in the frequency domain in order to remove high-frequency noise due to electronic ringing of the detector⁷. The measured memory efficiencies were fitted with an exponential decay function, yielding a mean memory lifetime of $1.7 \pm 0.1 \mu\text{s}$. This value agrees with the prediction of the decoherence model presented in Section 3.2.2, which takes into account magnetic dephasing of the Raman coherence due to a static residual from the earth’s magnetic field^[226].




Figure 4.3: Process fidelities (red squares) and memory efficiencies (blue dots) as a function of storage time. The efficiencies are fitted with an exponential decay function.

The sharp decrease of memory efficiency with storage time – to approximately 20% of its maximum value for retrieval after $2.5 \mu\text{s}$ – is in marked contrast to the behaviour of the process fidelity. While the fidelity was highest for a storage time of 12.5 ns ($F = 97 \pm 1\%$), it remained consistently above 84% for storage times of up to $1.5 \mu\text{s}$ ⁸. This discrepancy suggests that the coherence of the memory process remains unaffected by the decay of efficiency with increasing storage time.

One can therefore conclude that the decoherence mechanisms, which cause loss of signal

⁷Data processing was carried out by Patrick Michelberger.

⁸At longer storage times, the retrieved signal was not sufficiently bright to allow a meaningful reconstruction of the process matrix as a result of reduced memory efficiencies.

amplitude, do not reduce the phase coherence of the retrieved signal. This result implies that, even at low memory efficiencies, the fidelity of the stored information is not degraded, demonstrating the robustness of the polarisation coherence against memory loss.

As mentioned, the dominant limitation on memory lifetime is magnetic dephasing of the spin coherence due to stray magnetic fields. This seemingly calls into question the above observation that the process fidelity remains approximately constant at longer storage times. However, this apparent contradiction can be resolved by considering the orthogonality condition of the Raman interaction (see Section 4.1.1); during retrieval, the control pulse selects only the component of the spin wave which remains coherent with the input. This ensures that the retrieved field has the same polarisation as the original signal so that the polarisation state, and therefore the process fidelity, are maintained.

Chapter 5

Single photon storage

The storage and retrieval of single photons can be regarded as one of the principal objectives which has incentivised and guided our investigation of the Raman memory so far. The demonstration of true quantum operation is indeed a critical step towards implementing the Raman memory in large-scale temporal multiplexing applications, thus facilitating scalable photonic networks.

Remarkable advances have been made throughout the field of quantum memories, for instance single-photon storage has been demonstrated in cold atomic ensembles^{[113] [162] [264]} and via pre-programmed AFC echoes in cryogenic rare-earth-doped crystals^{[265] [266]}. Furthermore, narrowband storage and retrieval of non-classical states of light has been realised in room-temperature quantum memories^{[157] [159] [267]}. However, a scalable solution combining the key requisites of high-bandwidth storage, on-demand retrieval as well as low-noise room-temperature operation remains elusive.

The following chapter details our recent work concerned with interfacing the Raman memory with a source of broadband, heralded single photons generated via spontaneous parametric downconversion. As we shall see, the sole remaining challenge is the suppression

of noise during the memory interaction which currently obfuscates the non-classical photon statistics of the retrieved field. The second part of this chapter focuses on the characteristics of the underlying noise mechanism and addresses the main strategies for noise suppression.

The experimental work presented in this chapter is the result of close collaboration between myself and fellow PhD student Patrick Michelberger. The complexity and scope of the single-photon memory experiment meant that over the course of more than one year many parts of the experimental set-up were modified and measurements repeated by either one of us. While I was principally responsible for re-building the Raman memory experiment and investigating the source of noise, Patrick Michelberger set up, tested and optimised the single-photon source. As these two lines of work converged, we worked together on interfacing the memory with single photons, implementing feed-forward control and conducting a thorough characterisation of the memory performance.

5.1 Single photon source

As noted before, one of the distinguishing features of the Raman memory is its particularly broad acceptance bandwidth exceeding 1 GHz. This property, which is coupled with room-temperature operation, makes it possible to interface the memory directly with a standard downconversion source of single photons. In contrast, this type of source generally requires a high-finesse optical cavity in order to be compatible with narrow-band memories^[268]. We will now turn to a description of the single-photon source used in the following experiments.

5.1.1 SHG and SPDC

Single photons were generated by near-degenerate type-II spontaneous parametric downconversion (SPDC) in a periodically-poled potassium-titanyl-phosphate (PPKTP) waveguide,

pictured in Figure 5.1. During type-II SPDC, a pump photon is converted into two orthogonally polarised signal and idler (herald) photons whose frequencies and wave vectors satisfy both energy and momentum conservation¹. In bulk non-linear $\chi^{(2)}$ -crystals, these correlated photon pairs are emitted along a set of trajectories that satisfy the aforementioned boundary conditions. In contrast, the waveguide configuration adds a further boundary condition and confines photon pair emission to collinear, well-defined spatial modes^[114]. Compared to SPDC in bulk crystals, the phase matching condition is modified by the discrete transverse modes that are supported by the waveguide structure^{[270] [271]}.




Figure 5.1: (a) Input facet of the PPKTP waveguide chip. (b) Engraved waveguides of widths 4 - 2 μm (left to right). (c) EM-CCD image showing the spatial mode of the SPDC photons from the 3 μm waveguide.

The transverse confinement of the pump beam as well as the long interaction length result in improved pair generation efficiencies without the need for higher pump intensities^[269]. This is a particularly important advantage for the Raman memory experiment, since a less stringent requirement for the pump power means that more power is available for the memory control pulses. In addition, SPDC emission into collinear spatial modes can improve photon collection efficiencies and therefore increase the attainable heralding efficiencies.

The single-pass waveguide was pumped with 426-nm pulses, derived from 852-nm Ti:Sa pulses which were frequency-doubled in a 2-mm bulk PPKTP crystal via second harmonic generation (SHG). The SHG conversion efficiency was measured to be approxi-

¹Periodic poling of the KTP chip results in the quasi-phase-matching condition $\vec{k}_p = \vec{k}_s + \vec{k}_i + \vec{K}$, where $|\vec{K}| = 2\pi/\Lambda$ and $\Lambda = 10 \mu\text{m}$ is the poling period^[269].

ately $\eta_{\text{SHG}} \approx 0.005\%/W$. As illustrated in Figure 5.2, the SHG radiation was separated from the remaining un-converted pulses using a dichroic mirror. These un-converted pulses were passed through the Pockels cell, providing the control pulses for storage and retrieval, whereas the SHG field was fibre-coupled and directed towards the PPKTP waveguide (*AdvR*). A $40\times$ objective was used to focus the pump field into the waveguide which has a cross-section of $3\ \mu\text{m} \times 6\ \mu\text{m}$ and is enscribed in a 2-cm PPKTP sample. SPDC was implemented with an average input power of approximately 1 mW and a waveguide coupling efficiency of $\eta_{\text{WG}} \approx 10\%$.

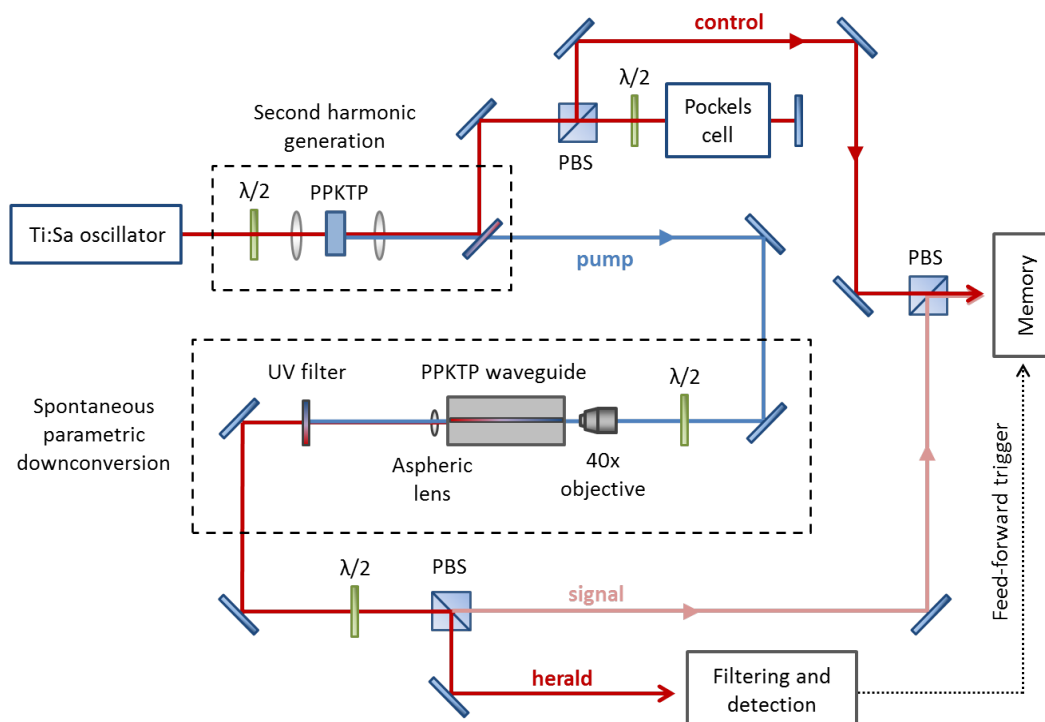


Figure 5.2: Set-up of the single photon source. The pump field is produced by frequency-doubling the Ti:Sa laser through second harmonic generation in a bulk PPKTP crystal. Type-II spontaneous parametric downconversion inside a PPKTP waveguide gives rise to orthogonally polarised signal and idler photons which are separated on a PBS.

In order to ensure stable SHG and SPDC efficiencies, the temperature of both the PPKTP bulk crystal and the waveguide were actively stabilised using separate temperature

controllers (*Newport*, Model 235). The optimal temperature for SHG and SPDC was determined to be 43.6°C and 27.7°C, respectively.

The SPDC-generated signal and idler photons were then UV-filtered to remove the residual pump field transmitted through the waveguide. Since the signal and idler photons were orthogonally polarised, they could be separated on a PBS and coupled into single-mode fibres with efficiencies of $\eta_{\text{signal}} = 72 \pm 2\%$ and $\eta_{\text{herald}} = 73 \pm 2\%$, respectively. The fibre-coupled herald was directed to a frequency-filtering stage, described in the following section, while the signal was interfaced directly with the memory set-up.

5.1.2 Herald filtering and detection

Due to the broad phase-matching bandwidth of the SPDC source, photon pairs were produced with a bandwidth of approximately 1 THz, i.e. three orders of magnitude greater than the acceptance bandwidth of the memory. However, by spectrally filtering the herald photon, the marginal spectrum of the signal photon could be projected into two-photon Raman resonance and spectrally matched to the bandwidth of the control field. This was achieved with a 0.01 nm-bandwidth volume holographic grating filter (*ONDAX*), and three air-spaced Fabry-Perot etalons, each with a bandwidth of ~ 1.5 GHz, as shown in Figure 5.3. Two of these etalons had a free-spectral range (FSR) of 18.2 GHz, while the third etalon, used in double-pass, had an FSR of 103 GHz.

Energy conservation during the SPDC interaction requires that the frequencies of the signal and herald photons satisfy the condition $\omega_{\text{sig}} + \omega_{\text{her}} = \omega_{\text{pump}}$. Hence, in order for the signal photon to be detuned by $\Delta - \Delta_{\text{HFS}} = (15.2 - 9.2)$ GHz = 6 GHz from the $6^2S_{\frac{1}{2}}F = 3 \rightarrow 6^2P_{\frac{3}{2}}$ transition, the detuning of the herald frequency had to equal $\Delta_{\text{her}} = 24.4$ GHz. Conditional on detection of a frequency-filtered herald photon, memory operation

would therefore only be attempted with spectrally-matched signal photons. The herald filter set-up was tuned to resonance at Δ_{her} by locking the Ti:Sa laser to the herald frequency and aligning the filters with a bright alignment beam, fibre-coupled in front of the Pockels cell. The peak transmission of the combined frequency filters was measured to be $T = 31 \pm 2\%$. The filtered herald photon was coupled into a multi-mode fibre and detected with a *Perkin Elmer* single-photon counting module (SPCM).




Figure 5.3: Filtering and detection set-up for the SPDC herald photon, using a holographic grating filter and three Fabry-Perot etalons. Detection events measured with a single-photon counting module (SPCM) are used to trigger the memory experiment, as discussed in Section 5.2.2.

5.2 Interfacing the memory with SPDC photons

5.2.1 Re-building the Raman memory

The previous Raman memory experiments, including polarisation storage, were conducted in a laboratory which – at the time – lacked the necessary temperature stability and did not offer sufficient space for the SPDC set-up. I therefore re-built the experiment in a more suitable laboratory with assistance from Duncan England. This also offered the opportunity for technical improvements to the experiment, the most important of which I will describe in the following sections.

Beam pointing stabilisation

Due to the relatively long free-space distances between optical components (up to several metres), the Raman experiment is highly sensitive to changes in the pointing of the Ti:Sa laser beam. These changes in beam position and direction may be caused by mechanical vibrations, thermal drifts or accidental misalignment. Moreover, slight perturbations of the Ti:Sa pointing occurred whenever the laser frequency was changed.

In order to counteract the effects of these disturbances and expedite alignment of the experiment, active stabilisation of the beam pointing was implemented. This was achieved by measuring changes in the direction of the beam path prior to the Pockels cell using CMOS cameras (*Mightex* BCE-B013-U) positioned behind two mirrors in the set-up. The mirrors were separated by ~ 2 m to increase the sensitivity with which changes in the beam pointing could be detected. A weak pick-off transmitted through each of the two mirrors was focused onto the camera so that the position of the beam centroid on the camera's active area could be determined. Changes of the centroid position relative to a fixed setpoint were monitored in both transverse directions. This provided the error signal for PID feedback using a commercial locking software (*KMLabs*, PointLock 2.0). The PID output was connected to a Picomotor Ethernet controller (*New Focus*, Model 8752) which adjusted the tip and tilt degrees of freedom of two motorised mirror mounts (*New Focus*, Model 8807) in the set-up. These motorised mounts were controlled via Picomotor actuators on two of its adjustment axes, allowing alignment with an angular resolution of $0.7 \mu\text{rad}$. Motorised control of the beam pointing had the additional benefit of providing an audible alert, whenever a sudden perturbation of the alignment caused the mirror mounts to be re-adjusted, accompanied by a distinct whirring sound.

This stabilisation method was tested by monitoring the coupling efficiency of the Ti:Sa

laser into a single-mode fibre (SMF) at a position downstream from the motorised mirrors, while purposefully disturbing the beam pointing. The feedback control readily compensated for these disturbances and repeatably recovered the SMF coupling efficiency. Active beam pointing stabilisation therefore greatly improved the robustness of the experimental set-up against misalignment.

Noise reduction

Previously, the unconditional noise floor of the Raman memory was measured to be approximately 0.25 photons per pulse (see Section 3.2.2). This noise floor included contributions at the Stokes (signal) frequency ν_s and at the anti-Stokes frequency $\nu_a = \nu_s + 2\Delta_{\text{HFS}}$, the origin of which will be discussed in Section 5.4. Both of these frequencies were allowed to pass the filtering stage, as the free-spectral range of the Fabry-Perot etalons was approximately equal to the frequency difference $\nu_a - \nu_s \approx 18$ GHz. By adding a fourth etalon with an FSR of 103 GHz, the anti-Stokes frequency component could be attenuated, as described in Section 3.1.6.

The amount of noise suppression offered by this new filtering set-up was tested by recording the detected noise count rates, while heating the vapour cell from 35°C to an operating temperature of 80°C. Over this temperature range, the vapour pressure of caesium is expected to change by a factor of approximately 40, according to the empirical vapour-pressure model for the liquid phase of caesium², given by

$$\log_{10} P_v = 2.881 + 4.165 - \frac{3830}{T} \quad (5.1)$$

where $T > T_M$ is the temperature in Kelvin and P_v is the pressure in Torr^{[229] [272]}. Since

²The melting point of caesium is $T_M = 28.5^\circ\text{C}$ ^[229].

the atomic number density is proportional to the vapour pressure for an ideal gas, one would therefore expect a significant drop in noise counts at low temperatures. However at temperatures of $T < 45^\circ\text{C}$, the noise appeared to plateau at approximately 30% of its high-temperature level, even though the number of gaseous atoms had gone down by 98% in the same temperature interval (see Figure 5.4). This strongly suggests that the detected noise contains an atom-independent contribution, most likely due to leakage of the control field through the filtering stage. However, by double-passing the 103-GHz etalon, the noise counts at low temperatures were reduced to below detectable levels, as shown in Figure 5.4. We can therefore conclude that the double-pass configuration, which provided an additional attenuation of 20 dB, successfully removed residual control photons.




Figure 5.4: Noise counts per 10-second time interval in the read-in and read-out time bins measured as a function of temperature for single-pass (blue/red) and double-pass (green/orange) etalon configurations. The solid line is the predicted caesium vapour pressure based on the empirical model given in reference^[229]. Error bars account for Poissonian errors in the count rates.

As a result of improved frequency filtering of both the anti-Stokes and control fields, the unconditional noise floor was significantly reduced from an average of 0.25 to 0.15 ± 0.05 per pulse in the read-out time bin. While this is an important advance towards near-noiseless Raman memory operation, we will see that further improvement is critical for the storage of single photons.

5.2.2 Feed-forward operation of the memory

One of the key requirements for memories used in temporal multiplexing applications is the ability to retrieve a stored excitation on demand. In the case of coherent-state storage, the signal was generated using an EOM, triggered by the intracavity photodiode signal of the Ti:Sa oscillator, as described in Section 3.1.4. Memory operation, including on-demand retrieval, could therefore be prompted using this same trigger. In contrast to coherent-state input signals, which are incident on the memory at regular intervals, the SPDC input photons are generated probabilistically. Hence, it was necessary to operate the memory in a feed-forward configuration, whereby detection of a herald photon would trigger the preparation of the control pulses for storage and retrieval. This was achieved by triggering a digital delay generator (DG535) with the herald photon-detector. Suitably delayed gating pulses from the DG535 were then used to trigger the Pockels cell so that pulse-picking of the control field was made conditional on a herald detection event (Figure 5.5).




Figure 5.5: Timing diagram of feed-forward operation: Detection of a herald photon triggers the Pockels cell via a digital delay generator (DG535) such that control pulses temporally overlap with SPDC signal photons.

For storage times exceeding a few tens of nanoseconds it is essential to turn off the optical pumping beam, since re-initialisation of the caesium atoms would destroy any spin wave coherence. However, the AOM used to switch the pumping beam suffers from a particularly long response time on the order of $1 \mu\text{s}$ which exceeds the available fibre delay of the heralded

SPDC signal photons. For this reason, the optical pumping beam was continuously present during SPDC storage. All memory experiments were therefore conducted for a storage time of 12.5 ns, corresponding to the separation of consecutive control pulses.

This limitation does not arise for coherent-state storage which is triggered deterministically, rather than from probabilistically occurring herald detection. With a deterministic trigger, the AOM switch-off could be delayed by one repetition period so that the pump was turned off for the following memory cycle. In particular, an experimental repetition rate of 5.72 kHz corresponds to 175 μ s between successive storage events which suffices for AOM switching.

5.2.3 Data acquisition

The signal emerging from the memory was both frequency- and polarisation-filtered using the same set-up as in previous experiments (see Section 3.1.6) and detected with *Perkin Elmer* SPCMs. Detection events were counted by a field-programmable gate array (FPGA) in coincidence with the herald SPCM³. In addition, a time-to-amplitude converter (TAC) and multi-channel-analyser (MCA) were used to record arrival time histograms between signal and herald photons. The output of the herald and signal SPCMs were connected to the TAC in order to provide two distinct TTL pulse triggers. Proportional to the temporal delay between these two triggers, the TAC produced a voltage signal ranging from 1 - 10 V. The MCA was then used to convert this range of voltages into corresponding time bins, thereby creating a histogram of arrival times. The number of time-bins was set by selecting a time interval (1 - 10000 ns) and resolution (up to 16384) which together defined the temporal width of each bin. For the current experiment, the bin size was set to \sim 30 ps.

³PhD student Justin Spring provided valuable assistance with the FPGA programming.

5.2.4 Temporal and spatial overlap

For efficient Raman coupling, the signal and control pulses have to be overlapped both temporally and spatially inside the memory medium. Implementing feed-forward operation, as described in Section 5.2.2, requires that the SPDC signal photon is suitably delayed to compensate for the electronic delays of the herald detector and the Pockels cell. Coupling the signal photon into an 83-m single-mode fibre delay line ensured approximate temporal overlap between the signal and the control pulses. Since the overlap between the coherent-state signal and the control was already optimised with respect to memory efficiency, the remaining task was to fine-tune the overlap between the SPDC photon and the classical signal pulse. This was achieved by recording arrival time histograms of both signals using a TAC (see Section 5.2.3 above). The overlap between these histograms was optimised by adjusting a delay stage in the coherent-state signal path. The delay of the control pulse was then adjusted accordingly.

The fibre-coupled SPDC and coherent-state signals were combined on a PBS and coupled into another SM-fibre with a length of 8 m. The output of this fibre was then spatially overlapped with the control pulse on a second PBS, as shown in Figure 5.6. The control itself was fibre-coupled in order to achieve a clean Gaussian mode profile and improve the spatial overlap within the memory.

By rotating a $\lambda/2$ -waveplate in front of the signal input port of the PBS used to recombine the signal and control fields, it was possible to switch quickly between coherent-state and SPDC storage. This was particularly convenient, because the memory efficiency could be optimised more readily using coherent states as input signals.

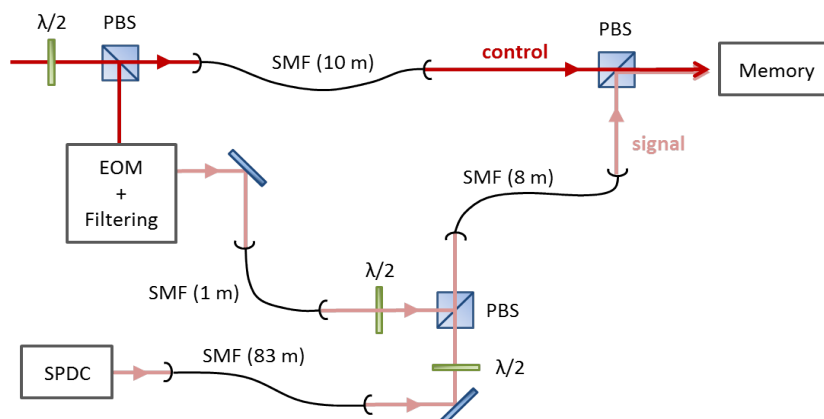


Figure 5.6: Spatial and temporal overlap of the control field with the coherent-state and SPDC input signals.

5.3 Memory operation

Having interfaced the memory set-up with the SPDC source, we are now able to conduct a thorough investigation of the memory performance. As we shall see below, this will allow us to demonstrate for the first time the storage of GHz-bandwidth heralded single photons in a room-temperature Raman memory. We will begin by considering the preparation efficiency of the memory input, i.e. the heralding efficiency of the SPDC source, before examining the total memory efficiency and the signal-to-noise ratio of the retrieved state.

5.3.1 Heralding efficiency

We will define the heralding efficiency η_{herald} as the probability with which a SPDC signal photon is incident at the input of the memory, given the detection of a filtered herald photon. Signal photons transmitted through the memory and the filtering set-up (see Section 3.1.6) were split on a PBS and measured with two *Perkin Elmer* SPCMs placed at either output port. As we shall see in Section 5.3.4, this set-up will be used to perform second-order autocorrelation measurements.

In order to determine η_{herald} , coincidences between the herald and both signal detectors (labelled H and V) were recorded, whilst optically pumping the atoms and blocking the control field. These coincidences, given by $c_{\text{her,sig}} = c_{\text{her,H}} + c_{\text{her,V}}$, were normalised by the number of herald counts c_{her} . This ratio is then corrected for the total transmission loss of the signal photon through the vapour cell and filter stage, $T_{\text{tot}} \approx 10\%$, as well as the detector efficiency $\eta_{\text{SPCM}} \approx 50\%$. The heralding efficiency is then given by

$$\eta_{\text{herald}} = \frac{c_{\text{her,sig}}}{c_{\text{her}} \cdot T_{\text{tot}} \cdot \eta_{\text{SPCM}}}. \quad (5.2)$$

Here, we have ignored contributions to the count rates from higher-order SPDC emissions, since the SPDC source was operated in a regime of low pump powers. Hence, all but single-pair emissions could be neglected. This approximation is substantiated by the low heralded $g^{(2)}$ -autocorrelation⁴ of the SPDC input photons at zero time delay which was measured to be $g^{(2)}(\tau = 0) = 0.016 \pm 0.004$.

After carefully optimising the SPDC spatial mode and minimising transmission losses through the signal path, a heralding efficiency of $\eta_{\text{herald}} = 0.22$ was achieved. This value is limited by scattering losses within the waveguide as well as the coupling efficiency between the waveguide output and the memory cell⁵.

5.3.2 Memory efficiency

The memory efficiency was determined by recording coincidence count rates between the signal and herald detectors in the read-in and read-out time bins, $c_{\text{her,H/V}}^{\text{in/out}}$, respectively. In order to access all contributions to the detected signal, the coincidence rates were measured

⁴As derived in Appendix A, the $g^{(2)}$ -autocorrelation for a Fock state $|n\rangle$ is given by $g^{(2)}(\tau = 0) = 1 - 1/n$.

⁵The total transmission loss of the SPDC signal photons from the output facet of the waveguide to the front window of the vapour cell is estimated to be $\sim 50\%$.

for a set of four different combinations which are labelled according to the laser beams present during data acquisition. These are as follows:

1. *scd*: The signal (*s*), control (*c*) and optical pumping diode (*d*) beams are all present, enabling memory operation.
2. *sd*: The input signal is transmitted through the memory without storage, in the absence of the control field.
3. *cd*: The noise contribution is isolated by blocking the input signal.
4. *d*: The optical pumping background is measured in the absence of both signal and control fields.

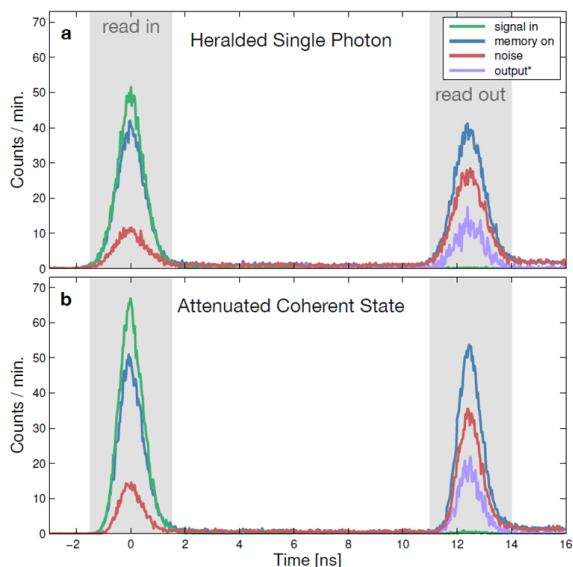


Figure 5.7: Storage and retrieval of a single photon (a), heralded with an efficiency of $\eta_{\text{her}} = N_{\text{in}} = 0.22$, and an attenuated coherent state (b) with an average photon number per pulse of $N_{\text{in}}^{\text{coh}} = 0.23$.

Arrival time histograms recorded with the TAC-MCA system are plotted in Figure 5.7. The input signal (*sd*) transmitted through the memory when the control field was blocked, is shown in green, while red traces correspond to noise emission in the read-in and read-out time bins (*cd*). Blue traces show the transmitted and retrieved signal during memory

operation (scd), including noise contributions. Finally, subtracting the noise present in the read-out time bin from the blue traces yields the effective retrieved signal, shown in lilac.

The coincidence rates $c_l^{\text{in/out}}$ with $l \in \{scd, sd, cd, d\}$ were obtained by integrating these traces and normalising by the sum of herald trigger events c_{her} . The total memory efficiency for storage and retrieval is then given by

$$\eta_{\text{tot}} = \frac{c_{\text{scd}}^{\text{out}} - c_{\text{cd}}^{\text{out}} - (c_{\text{sd}}^{\text{out}} - c_{\text{d}}^{\text{out}})}{c_{\text{sd}}^{\text{in}} - c_{\text{d}}^{\text{in}}}. \quad (5.3)$$

The uncertainty in η_{tot} was determined from Gaussian error propagation of the Poissonian errors on the count rates $c_l^{\text{in/out}}$, i.e. $\Delta c_l^{\text{in/out}} = \sqrt{c_l^{\text{in/out}}}$.

For SPDC input photons, heralded with an efficiency of $\eta_{\text{her}} = N_{\text{in}} = 0.22$, the memory efficiency was found to be $\eta_{\text{tot}} = 21.1 \pm 1.9\%$. This was benchmarked against the memory efficiency obtained for weak coherent states attenuated to an average photon number per pulse of $N_{\text{in}}^{\text{coh}} = 0.23 \approx \eta_{\text{herald}}$. With these input states, the memory efficiency amounted to $\eta_{\text{tot}} = 29.0 \pm 0.9\%$. The difference in efficiency can be ascribed to a residual spectral mode mismatch between the heralded SPDC photons and the control pulses. This reduces the efficiency with which SPDC photons can be stored compared to coherent states whose spectra are automatically mode-matched with the control. In contrast, the spectral mode of the SPDC signal is set by the spectral filtering of the herald photon so that small discrepancies can arise between signal and control modes.

5.3.3 Signal-to-noise ratio

The detected noise (cd) contained on average $\epsilon_{\text{in}} = (6 \pm 2) \cdot 10^{-2}$ photons per pulse in the read-in time bin and $\epsilon_{\text{out}} = (15 \pm 5) \cdot 10^{-2}$ photons per pulse during read-out⁶. The resulting signal-to-noise ratio (SNR), i.e. the ratio of retrieved signal photons to noise photons, is $\text{SNR} = \eta_{\text{tot}}\eta_{\text{herald}}/\epsilon_{\text{out}} = 0.3 \pm 0.1$. Since the SNR is directly proportional to η_{herald} , the achievable heralding efficiency sets tight tolerances on the unconditional noise floor of the memory.

Following our study of the memory efficiency and the signal-to-noise ratio, we can gain further insight into the memory performance by considering the effect of the memory interaction on the photon number statistics of the initial input state. For future applications, it is of critical importance that the non-classical statistics of a single-photon input signal are preserved during storage and retrieval. We will therefore extend our analysis by examining the second-order autocorrelation function at zero time delay, i.e. $g^{(2)}(\tau = 0)$. Note that from now on we will use the shorthand notation $g^{(2)}(0) := g^{(2)}$. As derived in Appendix A, $g^{(2)}(0) < 1$ is a signature of non-classical light exhibiting sub-Poissonian photon statistics^[273]. This inequality will serve as the benchmark for the memory's ability to preserve non-classicality. Moreover, $g^{(2)}$ -measurements of the noise present in the retrieved field are an important means by which unwanted contributions can be characterised.

The experimental procedure used to obtain $g^{(2)}$ -measurements and the subsequent data analysis will be described in the following section. The results of these measurements and their implications for single-photon storage in the Raman memory will be discussed in Section 5.3.5.

⁶These numbers were determined by dividing the detected count rate $c_{\text{cd}}^{\text{in/out}}$ by the transmission of the filtering stage and the photon-detector efficiency, i.e. $\epsilon_{\text{in/out}} = c_{\text{cd}}^{\text{in/out}}/(T_{\text{filter}}\eta_{\text{SPCM}})$.

5.3.4 $g^{(2)}$ -autocorrelation measurements

In order to determine the photon statistics of the retrieved field, heralded $g^{(2)}$ -autocorrelation measurements were conducted using a Hanbury Brown-Twiss set-up^[274]. The signal transmitted through the filtering stage was split into two spatial modes on a PBS, as shown in Figure 5.8. Each mode was coupled into a multi-mode fibre and connected to an SPCM detector. We will label these detectors according to the polarisation of the detected signal mode, i.e. D_H and D_V .

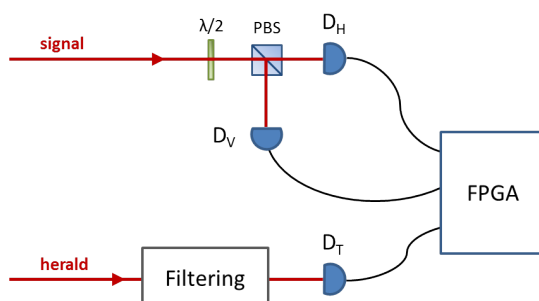


Figure 5.8: Detection set-up for $g^{(2)}$ -autocorrelation measurements. Coincidence events between the signal detectors D_H and D_V and the herald detector D_T are counted with the FPGA.

The normalised $g^{(2)}$ -autocorrelation, conditional on detection of a herald photon, is given by the ratio^[216]

$$g^{(2)} = \frac{p_{\text{her,H,V}}}{p_{\text{her,H}} \cdot p_{\text{her,V}}}, \quad (5.4)$$

where $p_{\text{her,H/V}}$ is the probability of coincidence detection between the herald detector D_T and one of the signal detectors $D_{H/V}$. The probability of all three SPCMs detecting a photon in coincidence is given by $p_{\text{her,H,V}}$. In terms of the detected coincidence counts, these probabilities were obtained from the ratios $p_{\text{her,H/V}} = c_{\text{her,H/V}}^{\text{in/out}}/c_{\text{her}}$ and $p_{\text{her,H,V}} = c_{\text{her,H,V}}^{\text{in/out}}/c_{\text{her}}$ for the input and output time bins. In this way, $g^{(2)}$ -measurements were taken using input signals derived from SPDC photons as well as coherent states with average photon numbers

ranging from $N_{\text{in}} = 0.23 \approx \eta_{\text{her}}$ to $N_{\text{in}} = 2.16$. In addition, the autocorrelation of the noise contribution was determined by blocking the input signal.

The experimental repetition rate of $f_{\text{rep}} = 5.3 - 7.3$ kHz was limited by the maximum repetition rate at which the Pockels cell could be driven without degrading the extinction ratio between picked and unpicked pulses. Since the Pockels cell was triggered upon detection of a herald photon, the repetition rate is $f_{\text{rep}} = c_{\text{her}}$, thus setting an upper bound on the brightness requirements of the SPDC source. In order to collect sufficient count rate statistics, it was therefore necessary to accumulate the coincidence data over many hours. In particular, the acquisition times needed for retrieved SPDC photons and coherent states of equivalent average photon numbers were approximately five hours each. Since the coincidence rates from noise counts were significantly lower, the required measurement time increased to 28 hours in order to determine the $g^{(2)}$ of the noise.

The total count rates for each experimental setting were summed over all measurement runs j , giving $c_{k,l}^t = \sum_{j=1}^{N_r} c_{j,k,l}^t$ where $k \in \{(\text{her}), (\text{her}, H), (\text{her}, V), (\text{her}, H, V)\}$ and N_r is the total number of runs for each setting $l \in \{scd, sd, cd\}$ and time bin $t \in \{\text{in}, \text{out}\}$. From these the coincidence probabilities $p_{k,l}^t = c_{k,l}^t / c_{\text{her},l}$, with $k \in \{(\text{her}, H), (\text{her}, V), (\text{her}, H, V)\}$, were determined. The $g^{(2)}$ was then evaluated by combining all measurement runs for a given setting and time bin, according to

$$g_{l,t}^{(2)} = \frac{P_{(\text{her},H,V),l}^t}{P_{(\text{her},H),l}^t \cdot P_{(\text{her},V),l}^t}. \quad (5.5)$$

The errors $\Delta g_{l,t}^{(2)}$ were obtained from Gaussian error propagation of the Poissonian errors $\Delta c_{k,l}^t = \sqrt{c_{k,l}^t}$ of the individual coincidence counts.

The long measurement times made it necessary to concatenate data obtained over mul-

multiple days. In order to test for comparable operating conditions, a set sequence of experimental runs was conducted each day prior to data collection. This sequence involved a 10-minute measurement of the $g_{sd}^{(2)}$ of the input signal, followed by two 30-minute measurements of $g_{scd}^{(2)}$ and $g_{cd}^{(2)}$.

The validity of this procedure can be checked by computing the $g_{j,l,t}^{(2)}$ -values for each measurement run j , from

$$g_{j,l,t}^{(2)} = \frac{c_{j,(\text{her,H,V}),l}^t \cdot c_{j,\text{her},l}}{c_{j,(\text{her,H}),l}^t \cdot c_{j,(\text{her,V}),l}^t}. \quad (5.6)$$

A double-sided, one-sample Student T-test was performed on the $g_{j,l,t}^{(2)}$ -values with $g_{l,t}^{(2)}$ from equation (5.5) as the assumed population mean. According to this test, the null hypothesis of $g_{l,t}^{(2)}$ being the mean of the j values $g_{j,l,t}^{(2)}$ cannot be rejected for any input type l and time bin t with a confidence of $\geq 95\%$. The validity of the T-test was then tested by performing a Shapiro-Wilk test on the $g_{j,l,t}^{(2)}$. This confirmed that the null hypothesis of the $g_{j,l,t}^{(2)}$ being normally distributed cannot be rejected with $\geq 95\%$ confidence for all l and t . Despite the fact that the coincidence counts $c_{j,k,l}^t$ are subject to Poissonian statistics, a normal distribution of $g_{j,l,t}^{(2)}$ is expected as a result of the central limit theorem⁷.

The corresponding memory efficiencies η_{tot} for SPDC and coherent state storage were determined using equation (5.3), averaged over all experimental runs. It should be noted that the count rates $c_l^{\text{in/out}}$ were scaled according to the measurement times for each setting l . Consistent memory operation was ensured by optimising the spatial overlap of the signal, control and pump beams between different measurement runs. Furthermore, the alignment of the signal filter stage and the SPDC heralding efficiency were monitored in order to avoid the effects of systematic drifts in the experimental apparatus.

⁷The statistical analysis was performed by Patrick Michelberger.

5.3.5 Results of the $g^{(2)}$ -measurements

The results of the $g^{(2)}$ -autocorrelation measurements are plotted in Figure 5.9, alongside the predictions of a theoretical model which will be discussed in Section 5.4.1.

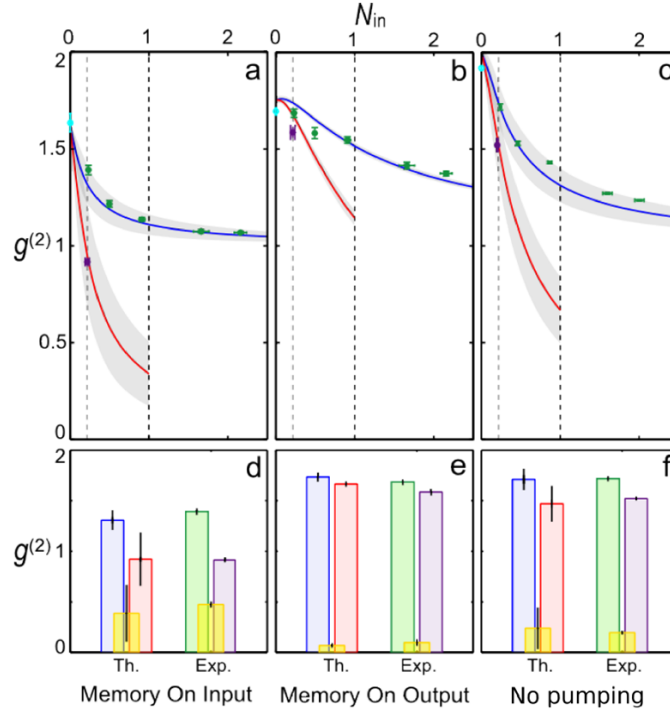


Figure 5.9: Measurements of the $g^{(2)}$ -autocorrelation for increasing average photon numbers per pulse, N_{in} , of the coherent-state input signal, for **a** signal transmitted through the memory without storage, **b** retrieved fields, **c** signal transmitted through the memory without optical pumping. Coherent-state input data is shown in green, purple data points represent a heralded SPDC input with $N_{\text{in}} = \eta_{\text{herald}}$ and cyan points are FWM noise. Errors bars derive from Poissonian counting statistics. The model predictions (see Section 5.4.1) for coherent-state and SPDC inputs are shown in blue and red, respectively. Shaded regions correspond to standard deviations determined from Monte-Carlo variation of the model parameters. Vertical grey and black dashed lines indicate $N_{\text{in}} = \eta_{\text{herald}}$ and $N_{\text{in}} = 1$ (perfect heralding efficiency), respectively. **d – f:** Bar plots comparing the theoretical model (Th.) and experimental data (Exp.) for $N_{\text{in}} \approx \eta_{\text{herald}}$, with colour coding as before. Yellow bars show the difference in $g^{(2)}$ -values of coherent-state and single-photon input signals.

In the absence of the control field, the $g^{(2)}$ for the SPDC input signal was found to be $g^{(2)} = 0.016 \pm 0.004$ which is close to the ideal value of $g^{(2)} = 0$, expected for single photons. For coherent input states, the $g^{(2)}$ -data obtained for all input photon numbers

N_{in} were combined by calculating the average over all individual $g_{j,l,t}^{(2)}$, weighted by the total measurement time for each N_{in} . This resulted in a weighted mean of $g^{(2)} = 1.01 \pm 0.01$, as expected for coherent input states. Appendix A contains a derivation of the $g^{(2)}$ -autocorrelation for coherent, thermal and single-photon states, each of which is governed by different photon statistics.

With the input signal blocked, we measured $g^{(2)} = 1.62 \pm 0.04$ and $g^{(2)} = 1.70 \pm 0.02$ for the noise in the input and output time bins, respectively. When the memory interaction was enabled, this noise contribution increased the $g^{(2)}$ of the transmitted and retrieved signal fields above the values measured without the control field present. In the read-in time bin, the heralded single photon displayed sub-Poissonian statistics just below the classical boundary with $g^{(2)} = 0.92 \pm 0.02$. In contrast, the $g^{(2)}$ of the coherent-state signal transmitted during read-in converged towards $g^{(2)} = 1$ with increasing input photon numbers N_{in} , as the SNR improved. In the read-out time bin, $g^{(2)} = 1.69 \pm 0.02$ for coherent input states with an average photon number comparable to the SPDC heralding efficiency, i.e. $N_{\text{in}} = 0.23 \approx \eta_{\text{her}}$. The statistics of the retrieved field are therefore indistinguishable from the noise. However, in the case of SPDC storage, the retrieved signal showed $g^{(2)} = 1.59 \pm 0.03$ which is a reduction by more than three standard deviations compared to the $g^{(2)}$ -values for coherent states and noise. This was determined by performing a one-sided, two-sample Welch test on $g_{j,scd,out}^{(2)}$ for both SPDC and coherent-state inputs. The null hypothesis, stating that these $g^{(2)}_j$ -samples have the same population mean, could be rejected with a confidence level of $\geq 99.7\%$, corresponding to a statistical significance of ≥ 3 standard deviations. Similar results were obtained when repeating the hypothesis test with the $g^{(2)}_j$ of the noise replacing those of the coherent-state signal. In order to corroborate this argument further, it was also shown that the population means of $g^{(2)}_j$ for coherent states and noise

could not be distinguished with any reasonable confidence.

The results of this analysis have therefore conclusively demonstrated a statistically significant difference between the $g^{(2)}$ for heralded SPDC photons and coherent-state signals. Although the $g^{(2)}$ of the former remained above the classical boundary due to the presence of noise, this nevertheless represents a clear signature of the non-classical input photon statistics on the retrieved field.

5.4 Modelling FWM noise

The results of the $g^{(2)}$ -autocorrelation described above impel us to examine rigorously the processes responsible for the measured photon statistics.

Previous studies of the Raman memory operated at the single-photon level have suggested that the dominant source of noise is four-wave mixing (FWM), seeded by spontaneous anti-Stokes scattering (see Section 3.2.2). This occurs when the control pulse addresses the initial ground state $|1\rangle$ into which the atoms have been optically pumped prior to storage. Absorption of a control photon ω_c results in the scattering of an anti-Stokes photon ω_a , detuned by $2\Delta_{\text{HFS}}$ from the signal frequency. This excites a spurious spin-wave coherence $|1\rangle\langle 3|$ which is uncorrelated with the storage of a signal photon and can act as the source of Stokes emission at the signal frequency ω_s . The schematic in Figure 5.10 illustrates this process.

The FWM interaction can be understood as a two-mode-squeezing operation which describes the creation of photon pairs in two separate modes a and b . The general Hamiltonian for two-mode squeezing is given by $H = \zeta^* \hat{a} \hat{b} + \zeta \hat{a}^\dagger \hat{b}^\dagger$, where ζ is a coupling parameter. By tracing over one of the modes, for instance through frequency-filtering, the photon statistics

of the other mode become identical to those of a thermal state^[275]. Hence, we should expect the $g^{(2)}$ -autocorrelation of the Stokes photon to yield $g^{(2)} = 2$, provided it is spatially and temporally single-mode (see Appendix A). The presence of multiple spatial modes on the other hand results in statistics which tend towards a Poissonian distribution⁸ with $g^{(2)} = 1$.




Figure 5.10: (a) Spontaneous anti-Stokes scattering at a frequency of $\omega_a = \omega_s + 2\Delta_{\text{HFS}}$ occurs when the control field couples to the initial state $|1\rangle$. (b) Double- Λ scheme showing the FWM process: Absorption of two control photons gives rise to the scattering of a non-degenerate Stokes and anti-Stokes photon pair at frequencies ω_s and ω_a , respectively.

In order to elucidate the FWM process and its effect on the output statistics of the retrieved field, a simple model was developed, based on the coherent off-resonant interaction between the incident light fields and the spin wave coherence^{[177][226][277]}. A brief overview of this model is given in the following section.

5.4.1 Coherent interaction model

The Maxwell-Bloch equations describing the Raman interaction in a one-dimensional propagation model were introduced in Section 2.2.2. We can take into account the effects of FWM noise by including the anti-Stokes field A , detuned from the excited state by an amount $\Delta' = \Delta + \Delta_{\text{HFS}}$, where $\Delta = 15.2$ GHz. The detuning of the anti-Stokes field

⁸For a multimode squeezed state it has been shown that $g^{(2)} \approx 1 + 1/K$ where K is the effective number of modes^[276]. This result was interpreted by noting that the thermal photon number distribution from each individual squeezer cannot be distinguished by the detector; the convolution of the second-order moments associated with each mode then gives rise to a Poissonian distribution in the limit of large K .

results in a modified Raman coupling constant $C' = \sqrt{Wd\gamma}/\Delta'$. The resonant optical depth d was previously measured to be ~ 1800 , provided the caesium vapour is kept at a constant temperature of 70°C . After adiabatic elimination of the excited state $|2\rangle$ ^{[208][278]}, the set of coupled partial differential equations is then given by

$$\begin{aligned}(\partial_z + i\delta k) S &= iCB \\ \partial_z A^\dagger &= -iC'B \\ (\partial_\omega + is) B &= iw (CS + C'A^\dagger),\end{aligned}\tag{5.7}$$

where δk is the phase mismatch between Stokes and anti-Stokes fields and we have defined the effective quasi-time-coordinate $\omega(\tau) = \frac{1}{W} \int_{-\infty}^{\tau} |\Omega(\tau')|^2 d\tau'$. $\Omega(\tau)$ is the Rabi frequency of the control pulse in terms of the retarded time $\tau = t - z/c$ with z running from 0 to $L = 7.5$ cm. The peak Rabi frequency can be estimated from the control pulse energy $\mathcal{E}_c \approx 10$ nJ, the pulse duration $T \approx 360$ ps and the beam waist of $300 \mu\text{m}$, giving $\Omega_{\text{max}} = 4.2$ GHz. The parameter W , which normalises the effective time $\omega(\infty) = 1$, is then equal to 3.2 GHz. Given a natural linewidth of $\gamma = 5.2$ MHz, we can determine the Raman coupling constants for the Stokes and anti-Stokes fields as $C = 0.82$ and $C' = 0.51$. The difference in population between the initial state $|1\rangle$ and the storage state $|3\rangle$ is given by $w = \langle |1\rangle \langle 1| \rangle - \langle |3\rangle \langle 3| \rangle = p_1 - p_3$. At room temperature, these two states are equally populated, i.e. $p_1 = p_3 = 0.5$ and $w = 0$. If we assume perfect optical pumping prior to each storage event, then we can set $p_3 = 0$ and $w = 1$. The wave vectors of the control, Stokes and anti-Stokes fields with centre frequencies ω_c , $\omega_s = \omega_c - \Delta_{\text{HFS}}$ and $\omega_a = \omega_c + \Delta_{\text{HFS}}$,

respectively, are

$$\begin{aligned}
k_{\Omega} &= \frac{L\omega_c}{c} + d\gamma \left(\frac{p_3}{\Delta} + \frac{p_1}{\Delta + \Delta_{\text{HFS}}} \right) \\
k_s &= \frac{L\omega_s}{c} + d\gamma \left(\frac{p_1}{\Delta} + \frac{p_3}{\Delta - \Delta_{\text{HFS}}} \right) \\
k_a &= \frac{L\omega_a}{c} + d\gamma \left(\frac{p_3}{\Delta + \Delta_{\text{HFS}}} + \frac{p_1}{\Delta + 2\Delta_{\text{HFS}}} \right).
\end{aligned} \tag{5.8}$$

The final term in each of these expressions accounts for dispersion due to the off-resonant interaction of the fields with the atomic transitions. The length of the cell L appears as a result of the normalisation of the z -coordinate. Dispersion between the Stokes and anti-Stokes fields leads to a FWM phase mismatch, given by $\delta k = 2k_{\Omega} - k_s - k_a$.

Finally, the strong control field perturbs the atomic energy levels by inducing an A.C. Stark shift which is quantified by the expression

$$s = \frac{W}{\Delta} + \frac{W}{\Delta'}. \tag{5.9}$$

The partial differential equations can be solved in terms of the system's Green's functions G_{jk} which represent the linear mapping between modes j and k . The signal S_{trans} transmitted through the memory during the storage process is then given by

$$S_{\text{trans}}(\omega) = \int_0^1 G_{SS}(\omega, \omega') S_{\text{in}}(\omega') d\omega' + \int_0^1 G_{AS}(\omega, \omega') A_{\text{vac}}^{\dagger}(\omega') d\omega' + \int_0^1 G_{BS}(\omega, z) B_{\text{therm}}(z) dz, \tag{5.10}$$

where A_{vac} and B_{therm} are the initial vacuum and thermal components of the anti-Stokes field and spin wave, respectively. The signal S_{ret} , retrieved after a storage time of T , can

be found using a similar expression,

$$S_{\text{ret}}(\omega) = \int_0^1 G_{SS}(\omega, \omega') S_{\text{vac}}(\omega') d\omega' + \int_0^1 G_{AS}(\omega, \omega') A_{\text{vac}}^\dagger(\omega') d\omega' + \int_0^1 G_{BS}(\omega, z) B_T(z) dz. \quad (5.11)$$

For storage times which are short compared to the memory lifetime, i.e. $T \ll O(\mu\text{s})$, we can neglect decoherence of the spin wave and set $B_T = B_{\text{out}}$ which, at the end of the interaction, is given by

$$B_{\text{out}}(z) = \int_0^1 G_{SB}(z, \omega) S_{\text{in}}(\omega) d\omega + \int_0^1 G_{AB}(z, \omega) A_{\text{vac}}^\dagger(\omega) d\omega + \int_0^1 G_{BB}(z, z') B_{\text{therm}}(z') dz'. \quad (5.12)$$

Using these solutions, we can determine the $g^{(2)}$ -autocorrelation according to

$$g_x^{(2)} = \frac{\int \int dt dt' \langle S_x^\dagger(t) S_x^\dagger(t') S_x(t') S_x(t) \rangle}{\left(\int dt'' \langle S_x^\dagger(t'') S_x(t'') \rangle \right)^2} = \frac{\int \int d\omega d\omega' \langle S_x^\dagger(\omega) S_x^\dagger(\omega') S_x(\omega') S_x(\omega) \rangle}{\left(\int d\omega'' \langle S_x^\dagger(\omega'') S_x(\omega'') \rangle \right)^2}, \quad (5.13)$$

where x represents ‘ret’ or ‘trans’. The expectation values of normally ordered products of the field operator S_x can be calculated from products of the Green’s functions^{[226][277]}. This follows from the bosonic commutation relations of the operators⁹ as well as the fact that the initial expectation values are known. The Green’s functions were computed using a previously-developed numerical code, although in principle analytical solutions can be found^[277].

The values of $g^{(2)}$ predicted by this calculation are plotted in Figure 5.9, alongside the experimental measurements. The width of the shaded error regions were determined by performing a Monte-Carlo error propagation. This involved computing the theoretical

⁹The spin wave operator satisfies a quasi-bosonic commutation relation, since $\langle [B_{\text{in}}(z), B_{\text{in}}^\dagger(z')] \rangle = \omega \delta(z - z')$.

prediction 1000 times using input parameters drawn from a Gaussian distribution with standard deviation set by the experimental uncertainties. These were estimated as $\sigma_{\text{std}}(d) = 100$, $\sigma_{\text{std}}(\gamma) = 0.5$ MHz, $\sigma_{\text{std}}(\Delta) = 200$ MHz, $\sigma_{\text{std}}(T_c) = 20$ ps and $\sigma_{\text{std}}(\Omega_{\text{max}}) = 100$ MHz.

The excellent agreement between theory and experiment confirms that the measured noise is indeed in accordance with FWM. Moreover, this FWM process can only be correctly modelled if the full coherent dynamics between the atomic system and the optical fields are taken into account. Building on our discussion in Chapter 2, the inclusion of the anti-Stokes field in the dynamical equations has allowed us to arrive at a more complete picture of the underlying physics. In addition to providing a description of the noise process, the coherent interaction model has therefore refined our understanding of the Raman memory interaction.

5.4.2 Incoherent interaction model

In order to provide further verification for the above description of FWM, we can test the measured data against predictions obtained from an incoherent interaction model. This model is based on the assumption that the individual photon statistics of the input signal $g_{\text{sig}}^{(2)}$ and the noise $g_{\text{noise}}^{(2)}$ can be combined incoherently, akin to the mixing of different modes on a beam splitter. The statistics of the multimode output, $g_{\text{tot}}^{(2)}$, are then simply given by a combination of $g_{\text{sig}}^{(2)}$ and $g_{\text{noise}}^{(2)}$.

In a recent paper, Elizabeth Goldschmidt demonstrated a method of reconstructing the mode distribution of an optical field containing multiple contributions, each with different classical or non-classical photon-number statistics^[279]. The probability distribution for each of these underlying modes can be described by a probability generating function $G(s) = \sum_n p(n)s^n$, where $p(n)$ is the probability that the mode contains n photons. For a two-

mode probability distribution this becomes $p(k) = \sum_n p_1(n)p_2(k-n)$. The generating function of the probability $p(k)$ is then given by the product of the generating functions of the component mode probabilities p_1 and p_2 . This may be interpreted in analogy with the convolution theorem for Fourier transforms. By extension, the generating function $G_{\text{tot}}(s)$ for an n -mode distribution is given by the product $G_1(s) \times G_2(s) \times \dots \times G_n(s)$ ^[280]. We can relate the probability generating function to the k^{th} -order intensity autocorrelation function at zero time delay via $g^{(k)}(0) = G^{(k)}(s=1)/(\mu_{\text{tot}})^k$, where $G^{(k)}(s=1)$ is the k^{th} derivative of $G(s)$ at $s=1$ and μ_{tot} is the mean photon number. It is then straightforward to show that the second-order autocorrelation function for a mixture of signal and noise modes is given by

$$g_{\text{tot}}^{(2)} = \frac{(N_{\text{sig}})^2 \cdot g_{\text{sig}}^{(2)} + 2N_{\text{sig}} \cdot N_{\text{noise}} + (N_{\text{noise}})^2 \cdot g_{\text{noise}}^{(2)}}{(N_{\text{sig}} + N_{\text{noise}})^2}, \quad (5.14)$$

where N_{sig} and N_{noise} are the number of signal and noise photons per pulse, respectively. When the memory is on, i.e. when the control field is unblocked, N_{sig} can be determined as follows: in the read-in time bin, N_{sig} is the fraction of the input signal transmitted through the memory, i.e. $N_{\text{sig}} = (1 - \eta_{\text{read-in}}) \cdot N_{\text{in}}^{\text{sd}}$, whereas in the read-out time bin $N_{\text{sig}} = \eta_{\text{tot}} \cdot N_{\text{in}}^{\text{sd}}$ is the retrieved field. For the case of heralded single photons, the read-in efficiency is equal to $\eta_{\text{read-in}} = 39.1\% \pm 3.5\%$, the total memory efficiency is $\eta_{\text{tot}} = 21.1\% \pm 1.9\%$ and the number of input photons is $N_{\text{in}}^{\text{sd}} = \eta_{\text{herald}} = 0.22 \pm 0.03$.

The predicted $g^{(2)}$ -values obtained using equation (5.14) are shown as dashed lines in Figure 5.11 for both the input and output time bins. This prediction consistently underestimates the experimental data, which allows us to reject the incoherent interaction model.

In evaluating equation (5.14) we have implicitly assumed that the photon statistics of the input signal are conserved during memory storage. One could argue that this is not

necessarily a valid assumption. In order to calculate $g_{\text{tot}}^{(2)}$, we would therefore have to replace $g_{\text{sig}}^{(2)}$ with modified photon statistics which are not directly accessible experimentally. However, this objection can be refuted by comparing the model prediction with measurements obtained when the optical pump beam was blocked. Having switched off the Raman interaction, $g_{\text{sig}}^{(2)}$ should remain unchanged. Moreover, N_{sig} is equal to the number of input photons without the need to take into account memory efficiency factors. Figure 5.11 shows that even in this case, in which all variables in equation (5.14) are directly measurable, the incoherent model fails to describe the observed statistics.

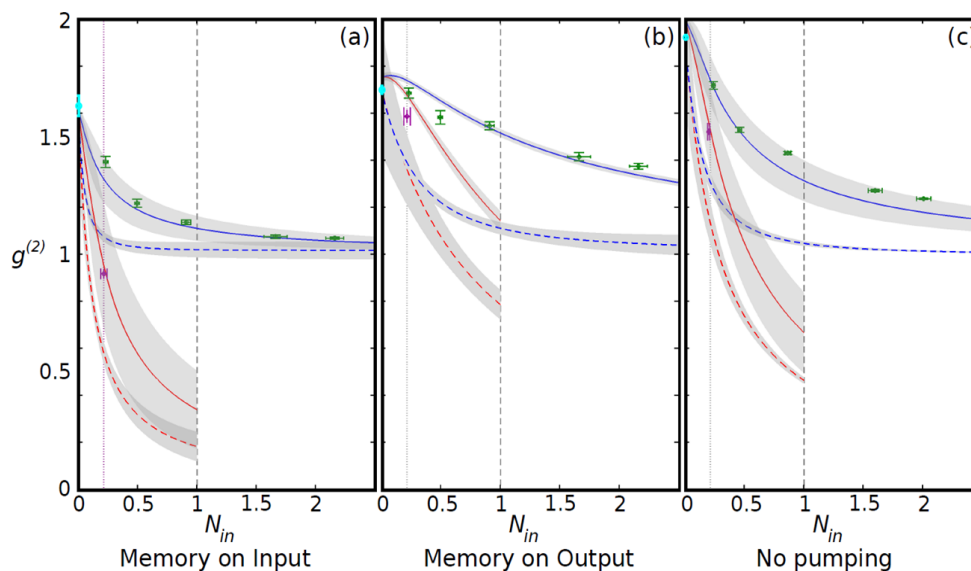


Figure 5.11: $g^{(2)}$ -predictions based on the incoherent interaction model in equation (5.14) (dashed lines), compared to experimental data (points) and predictions of the coherent $g^{(2)}$ -model (solid lines), for (a) fields transmitted in the read-in time bin during storage, (b) retrieved fields in the output time bin and (c) fields transmitted through the memory without optical pumping. Grey shaded areas denote errors derived from experimentally measured uncertainties on the variables in equation (5.14) and from propagating uncertainties in the experimental parameters through the FWM model with a Monte-Carlo simulation.

We can therefore conclude that the signal and noise modes cannot be treated as independent entities; instead the autocorrelation measurements reveal that the system can only be fully understood in terms of a coherent interaction model¹⁰.

¹⁰Due to the coherent nature of the FWM interaction, noise contributions cannot simply be subtracted

5.5 Further noise investigation

Following on from our theoretical model of FWM and in addition to the $g^{(2)}$ -autocorrelation measurements, I will now present further experimental evidence which supports our understanding of this noise process. As before, noise counts (control + optical pumping) were measured by blocking the signal field and recording detection events using the FPGA and TAC/MCA systems.

5.5.1 Build-up of FWM noise

A comparison of the FWM noise floor measured in the read-in and read-out time bins showed an increase by a factor of 2.5, from $\epsilon_{\text{in}} = (6 \pm 2) \cdot 10^{-2}$ to $\epsilon_{\text{in}} = (15 \pm 5) \cdot 10^{-2}$ (see Section 5.3.3). This is due to the fact that the noise in the read-out bin contained contributions from both the write and the read control pulses. Moreover, TAC traces recorded with a train of control pulses incident on the memory revealed a build-up of noise counts over successive time bins.

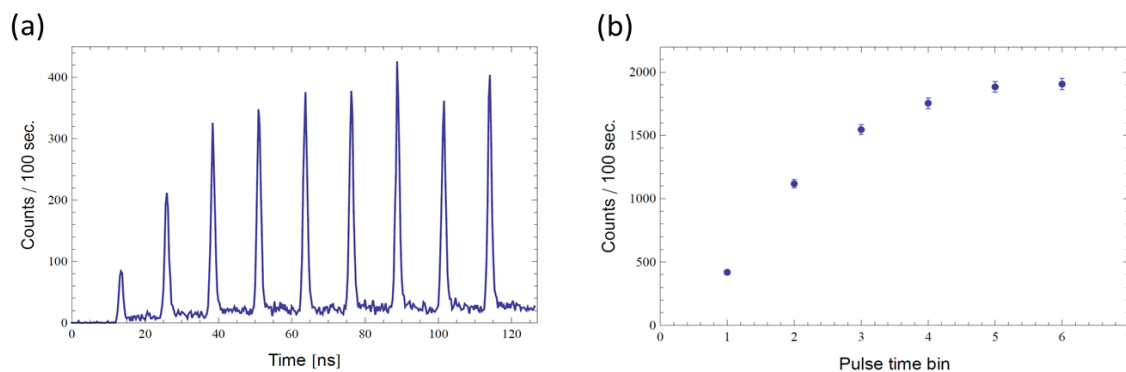


Figure 5.12: (a) TAC arrival time histogram for a train of nine consecutive control pulses. (b) Integrated pulse areas, showing build-up of noise counts in the first six time bins.

Figure 5.12 shows the integrated pulse areas of these TAC traces in each time bin. This demonstrates that in addition to instantaneous FWM noise, spurious spin wave components from the detected signal. For this reason, the memory efficiency was measured with bright coherent states for which the FWM noise was negligible.

excited by preceding control pulses can be *retrieved* across multiple bins. The rise-time of the noise build-up is expected to be correlated with the coupling strength of the FWM interaction, similar to the retrieval of a stored signal excitation with multiple control pulses (see Section 3.2.3).

5.5.2 Temperature dependence

According to the analysis presented in Section 5.4.1, the interaction strengths of the memory and FWM processes should have the same dependence on the optical depth d , since both coupling parameters C and C' are proportional to \sqrt{d} . To test this, the noise present during read-out was measured for different vapour cell temperatures, while recording the corresponding memory efficiency for bright coherent input signals on a fast photodiode. It should be noted that these measurements were taken with the improved frequency filtering in place (see Section 5.2.1) so that temperature-independent control leakage was negligible.




Figure 5.13: Memory efficiency (red) of bright coherent states and noise counts per 10-second time interval (blue) measured during read-out, as a function of vapour cell temperature. Below a temperature of 55°C, the memory efficiency became vanishingly small and the retrieved signal could not be distinguished from the background.

Based on the relation between optical depth and number density in equation (3.1), the temperature dependence both of the memory efficiency and the noise data is expected to

follow an exponential of the form given in the vapour-pressure model of equation (5.1). The close agreement between the two sets of data in Figure 5.13 clearly supports this assertion.

5.5.3 FWM coupling strength

Finally, the noise counts in the read-in and read-out time bins were measured for different control field powers P_c . From these, we can obtain an estimate of the energy \mathcal{E}_c in each control pulse by taking into account the experimental repetition rate, the number of control pulses per experimental cycle and the extinction of the Pockels cell¹¹. The results are plotted in Figure 5.14, alongside the corresponding memory efficiency, as measured with bright coherent states.




Figure 5.14: Memory efficiency (red) of bright coherent states and noise counts per 10-second time interval measured during read-in (blue) and read-out (green), for different values of the control pulse energy. The solid lines show quadratic fits of the data sets.

The Raman coupling constant, defined in terms of field amplitudes, is proportional to the square root of the control pulse energy \mathcal{E}_c , as shown in equation (2.44). Since the storage and retrieval efficiencies each scale approximately with the square of the coupling

¹¹The measured power P_c can be separated into two contributions, i.e. $P_c = \mathcal{E}_c N_c f_{\text{rep}} + \chi^{-1} \mathcal{E}_c \times (f_{\text{Ti:Sa}} - N_c f_{\text{rep}})$. The first term accounts for the $N_c = 18$ control pulses which are picked from the 80-MHz train of Ti:Sa pulses ($f_{\text{Ti:Sa}} = 80$ MHz) by the Pockels cell with a repetition rate of $f_{\text{rep}} = 4$ kHz and the second term accounts for the pulses that leak through the Pockels cell. Here, $\chi = 4000$ is the Pockels cell extinction, defined as the ratio between picked pulses and leakage.

constant, the total memory efficiency is proportional to the square of the control energy. This dependence holds in the regime of low control field powers to which we are currently limited experimentally. The efficiency data in Figure 5.14 was therefore fitted with a quadratic function. If the control field power were increased further, the memory efficiency would eventually saturate and then decrease as a result of the A.C. Stark shift (see equation (5.9)).

In contrast, noise originating from a FWM gain process is expected to scale exponentially with the control energy. This scaling can be understood by considering the effective two-mode squeezing operator $S_{s,a}(\zeta) = \exp\left[\frac{1}{2}(\zeta^* \hat{s} \hat{a} - \zeta \hat{s}^\dagger \hat{a}^\dagger)\right]$, where $\zeta = r \exp(i\theta)$ is the complex squeezing parameter^[281]. The magnitude r determines the interaction strength and is proportional to the intensity of the control field^[282]. The annihilation operators \hat{s} and \hat{a} refer to the Stokes and anti-Stokes modes, respectively.

Using the unitary transformation properties of the squeezing operator^[273], we can find an expression for the annihilation operator of the output field via the Bogoliubov transformation^[283],

$$\hat{s}_{\text{out}} = \hat{s}_{\text{in}} \cosh r + \hat{a}_{\text{in}}^\dagger e^{i\theta} \sinh r. \quad (5.15)$$

The number of excitations in the output mode is then given by

$$\langle \hat{s}_{\text{out}}^\dagger \hat{s}_{\text{out}} \rangle = \langle \hat{s}_{\text{in}}^\dagger \hat{s}_{\text{in}} \rangle \cosh^2 r + \langle \hat{a}_{\text{in}}^\dagger \hat{a}_{\text{in}} \rangle \sinh^2 r = \langle \hat{s}_{\text{in}}^\dagger \hat{s}_{\text{in}} \rangle \cosh^2 r + \sinh^2 r, \quad (5.16)$$

where, in the last step, we have used the bosonic commutator $[\hat{a}, \hat{a}^\dagger] = 1$ and assumed that $\langle \hat{a}_{\text{in}}^\dagger \hat{a}_{\text{in}} \rangle = 0$. Equation (5.16) demonstrates that both FWM contributions to the output mode scale exponentially with the strength of the squeezing parameter. The $\sinh^2 r$ -term in this equation corresponds to the spontaneous noise contribution which is a signature of any amplification process and independent of the input signal. Provided that the squeezing

parameter is sufficiently small, $\sinh^2 r$ can be approximated with a quadratic function. This approximation appears to be valid for our system, given that the dependence of the measured noise data with control pulse energy is well-described by a quadratic fit function (see Figure 5.14). The scaling of FWM noise is therefore expected to be similar to that of the overall memory efficiency for the range of available control field energies.

5.5.4 Spatial distribution of Stokes scattering

The spatial intensity correlation between Stokes and anti-Stokes fields has recently been studied experimentally through far-field imaging of the scattered radiation from warm rubidium vapour^[284].

Dabrowski *et al.* suggested that in certain applications it may be possible to exploit the directionality imposed by FWM phase matching for improved signal filtering^[284]. Since Stokes noise is mainly scattered into the spatial mode defined by the control field (in contrast to isotropically scattered spontaneous emission), a non-collinear configuration of signal and control beams may therefore allow some degree of spatial filtering of the noise. More specifically, angle-tuning the control with respect to the ensemble axis could potentially reduce the detected noise level in the fibre-coupled signal mode.

In order to test the viability of this option, a method had to be devised in order to pivot the control beam at the centre of the memory cell, so as to maintain spatial overlap with the signal field in the interaction region. For this purpose, the position and direction of the signal and control beams were monitored using the set-up of cameras described in Section 3.1.5. Separating the two beams by a small amount ε (on the order of a few millimetres) on one camera, while maintaining overlap on the other, made it possible to introduce a variable angle between the signal and control, as shown in Figure 5.15.

However, preliminary measurements suggested that any deviation from co-linearity quickly degraded the memory efficiency as a result of the reduction in the size of the interaction region. While further tests are needed in order to obtain conclusive results, it appears that a significant level of noise reduction cannot be achieved through this method of spatial filtering without sacrificing memory efficiency by a similar amount.




Figure 5.15: Angle-tuning the control field with respect to the signal beam axis by a small angle $\tan^{-1}(\varepsilon/d)$, where d is the path length difference between the two cameras.

5.6 Strategies for reducing FWM noise

5.6.1 Phase-matching condition

The FWM-interaction is subject to the conditions of both energy and momentum conservation. The former requires that the frequencies of the Stokes and anti-Stokes photons satisfy the equation $2\omega_c = \omega_a + \omega_s$, where ω_c is the frequency of the control field. Furthermore, momentum conservation imposes the phase-matching condition $2\vec{k}_c + \vec{k}_a + \vec{k}_s = 0$ ^[285]. This condition cannot be satisfied for collinear wavevectors due to dispersion between the Stokes and anti-Stokes fields which are detuned from the excited state by different amounts.



Figure 5.16: FWM phase matching condition for (a) collinear wavevectors in the absence of dispersion and (b) non-collinear Stokes and anti-Stokes wavevectors subject to dispersion.

As described in Section 5.4.1, the off-resonant interaction with an atomic transition modifies the wavevectors by an amount proportional to the detuning from that transition.

The explicit expression for the FWM phase mismatch is given by

$$\begin{aligned} \delta k &= d\gamma \left(\frac{2p_3 - p_1}{\Delta} + \frac{2p_1 - p_3}{\Delta + \Delta_{\text{HFS}}} - \frac{p_3}{\Delta - \Delta_{\text{HFS}}} - \frac{p_1}{\Delta + 2\Delta_{\text{HFS}}} \right) \\ &= d\gamma p_1 \left(\frac{2}{\Delta + \Delta_{\text{HFS}}} - \frac{1}{\Delta + 2\Delta_{\text{HFS}}} - \frac{1}{\Delta} \right), \end{aligned} \quad (5.17)$$

where we have assumed that energy conservation is satisfied and made the approximation $p_3 \approx 0$, i.e. perfect optical pumping. Phase-matched FWM with $\delta k = 0$ occurs if $\Delta_{\text{HFS}} = 0$, corresponding to a Λ -system with degenerate ground states. Conversely, the FWM-interaction would be suppressed if it were possible to operate the memory in a regime where $\Omega_c, \delta < \Delta < \Delta_{\text{HFS}}$, thus increasing the phase mismatch δk , while still satisfying the adiabatic approximation. By operating the memory closer to the atomic resonance, one would thereby increase the relative detuning of the anti-Stokes field. However, this potential route towards reducing the FWM noise contribution by increasing the amount of dispersion turns out to be unfeasible for our current system due to the onset of linear absorption for $\Delta < 6$ GHz. For future memory experiments, it may be possible to gain some advantage by using a vapour cell with lower buffer gas pressure in order to reduce the effects of collisional broadening on the resonance linewidth. However, some amount of buffer gas is nevertheless

essential, not only to increase the memory lifetime dramatically, but also to enable our current method of optical pumping with a narrow-band diode laser. This method relies on the fact that the velocities of the caesium atoms change randomly through collisions with the buffer gas so that their Doppler-shifted resonance frequencies *wander* across the entire Doppler profile. Eventually all atoms will drift into resonance with the narrow-band diode laser via this process of *spectral diffusion*. Optical pumping is therefore possible without the need to dither the frequency of the diode laser, provided the timescale of spectral diffusion is short compared to atomic diffusion out of the optical beam.

Alternatively, one could consider implementing the memory in different storage media which display larger Stokes shifts^[286]. This has already been drawn upon in a recent experiment demonstrating the low-noise storage of THz-bandwidth SPDC photons in the optical phonon modes of bulk diamond^[178]. By taking advantage of the 40-THz splitting between the ground and storage states as well as the high level of optical dispersion, Raman storage in diamond circumvents the problem of FWM noise. However, as noted earlier, storage times are unfortunately limited to a few picoseconds due to the rapid decay of optical phonons.

Upon inspection of equation (5.17), one could argue that an increase in optical depth would also increase the amount of phase mismatch, since $\delta k \propto d$. For instance, one may suggest simply using a longer vapour cell. Yet while FWM has a sinc^2 -dependence on the phase mismatch, the coupling strength of the FWM interaction, as a gain process, increases exponentially with d . This more than compensates for the greater dispersion so that, overall, FWM cannot be reduced by changing the optical depth.

It should be noted that we have so far neglected the noise contribution from spontaneous Stokes scattering as a result of imperfect optical pumping. Since this scattering process is not subject to a phase-matching requirement, there is no preferred direction of emission so

that scattering occurs isotropically over an angle of 4π . Compared to directional scattering processes, spontaneous Stokes scattering is suppressed by a factor of $4\pi/\Omega$, where Ω is the solid angle of detection. Since $\Omega \ll 4\pi$, it is therefore justified to neglect this source of noise. Moreover, previous studies have indicated that near-perfect optical pumping efficiencies can be achieved, given the available pump powers in our system^{[174][241]}.

5.6.2 Selection rules

A different approach to overcoming the common problem of FWM noise in warm-alkali-vapour memories was put forward by Walther *et al.* The proposed solution involves lifting the degeneracy of the Zeeman sublevels $|F, m_F\rangle$ and then optically pumping the atoms into the sublevel with the largest angular-momentum projection m_F along the optical axis. Figure 5.17 illustrates the two-photon Raman interaction which couples circularly polarised signal and control fields of opposite helicity. As a result of selection rules, the σ^+ -polarised control field cannot couple to the initial state $|1\rangle = |F = 4, m_F = 4\rangle$, since there exists no upper state $|F', m_{F'} = 5\rangle$ with an allowed transition. Hence, the control cannot generate a spurious spin wave coherence and FWM is suppressed.




Figure 5.17: Proposed scheme for suppressing FWM noise based on selection rules. Initially, the atoms are optically pumped into the extremal Zeeman state $|F = 4, m_F = 4\rangle$ (a) to which the σ^+ -polarised control field cannot couple (b), since there is no $m_F = 5$ state with an allowed transition. For clarity, the detuning Δ from the excited states is not shown.

Unfortunately, however, this potentially elegant solution is not viable for our system, since the matrix elements of the Raman processes involving two different intermediate states F' interfere destructively in the limit of large detunings. By calculating the Clebsch-Gordan coefficients of the initial, intermediate and final states, one finds that the Raman matrix elements for the two allowed transitions have equal amplitudes and opposite signs, i.e.

$$M_R = \langle F_{\text{final}} | \mathbf{d} | F'_1 \rangle \langle F'_1 | \mathbf{d} | F_{\text{initial}} \rangle = - \langle F_{\text{final}} | \mathbf{d} | F'_2 \rangle \langle F'_2 | \mathbf{d} | F_{\text{initial}} \rangle, \quad (5.18)$$

where $F_{\text{final}} = |F = 3, m_F = 2\rangle$ is the storage state $|3\rangle$, and F'_1 and F'_2 are the intermediate states $|F' = 3, m_{F'} = 3\rangle$ and $|F' = 4, m_{F'} = 3\rangle$ of the $6^2P_{3/2}$ manifold which are involved in the Raman interaction. The sum of these matrix elements cancel, thereby prohibiting memory operation.




Figure 5.18: Raman matrix elements associated with coupling via two intermediate states (here: $|F' = 3, m_{F'} = 3\rangle$ and $|F' = 4, m_{F'} = 3\rangle$) have equal amplitudes and opposite signs, resulting in destructive interference between these two interaction pathways.

The generality of this result for all alkali isotopes, as well as for both D_1 and D_2 transition lines, was recently demonstrated by Igor Vurgaftman and Mark Bashkansky^[287].

5.6.3 Density of states

A third strategy for reducing FWM noise originates from a density-of-states perspective. By reducing the density of states into which the anti-Stokes photon can be emitted, FWM is suppressed with respect to the memory process, resulting in an improvement of the signal-to-noise ratio.

The final part of this thesis is devoted to exploring this approach which was realised by operating the memory inside a low-finesse optical cavity. In the following chapter we will begin by theoretically examining the cavity strategy and its implications for the Raman memory interaction.

Chapter 6

The cavity Raman memory

A particularly promising strategy with which to address the key challenge of noise suppression involves a modification of the density of states by operating the memory inside an optical cavity. To this end, a cavity is chosen which can be tuned into resonance with the signal frequency, while simultaneously suppressing the anti-Stokes frequency.

In the first part of this chapter I will begin with a general theoretical description of the bare light-cavity system, before discussing the dynamics of the intracavity Raman memory interaction. The second half of this chapter is concerned with design considerations of the experimental cavity implementation, including the optimal cavity parameters, its geometry as well as methods of active stabilisation.

6.1 Input-output theory

Input-output (I-O) theory provides an important formalism with which to describe the coupling between external light fields and optical cavity modes. This approach was first developed in the mid 1980s by Collett and Gardiner^[288], Gardiner and Savage^[289], and Yurke^[290], at a time which also saw rapid advances in the formulation of cavity QED. I-O

theory yields a quantum-mechanical description of the dynamics governing the input, cavity and output fields of the linear light-cavity system. In the absence of intracavity losses, the unitarity of the evolution equations is ensured by verifying that the commutation relations agree with the fundamental equal-time field commutator^[291].

I-O theory has been commonly limited to the *good cavity* regime of high Q-factors^[288]^[292]. In this limit, the bandwidth of the optical field has to be significantly smaller than the cavity's free spectral range so that the effects of pulse evolution during a cavity round-trip can be ignored. However, in a recent paper, Raymer and McKinstrie extended I-O theory for arbitrary light-cavity coupling strengths^[293], based on previous work in references^[291]^[294]. This generalised approach pertains to the Raman cavity design (see Section 6.10.3) which necessitates operation in the low-Q-factor regime. The following is an overview of their results.

6.1.1 Field propagation equations in the space-time domain

Raymer and McKinstrie consider a travelling-wave ring cavity with a single optical input field at a junction with transmission and reflection coefficients \sqrt{R} and \sqrt{T} , respectively, which satisfy $R + T = 1$.




Figure 6.1: Ring cavity geometry, showing the input field $A(t)$, the output field $B(t)$ and the cavity field $C(t)$. External reflections at the input junction give rise to a phase flip, i.e. the associated amplitude reflection coefficient is negative.

Figure 6.1 illustrates the cavity geometry as well as the phase convention used in the following derivation, i.e. external reflections cause a phase flip. The intracavity field amplitude $C(z, t)$ evolves according to the travelling-wave Maxwell equation:

$$(\partial_t + v_g \partial_z)C(z, t) = \alpha P(z, t), \quad (6.1)$$

where v_g is the group velocity. The electric dipole polarisation $P(z, t)$ is zero for an empty cavity and α parametrises the coupling strength. The boundary condition at the coupling junction inside the cavity relates the cavity field and the input field $A(t)$,

$$C(0_+, t) = \sqrt{R}C(l_-, t) + \sqrt{T}A(t), \quad (6.2)$$

where 0_+ is the position inside the cavity right after the coupling junction and l_- is the position after one round-trip.

Similarly, the output field $B(t)$ must satisfy the boundary condition

$$B(t) = \sqrt{T}C(l_-, t) - \sqrt{R}A(t), \quad (6.3)$$

where the phase flip upon external reflection is expressed by the minus sign.

These boundary conditions at the coupling junction are simply the well-known beam splitter relations and can be conveniently combined in matrix form

$$\begin{pmatrix} B(t) \\ C(0_+, t) \end{pmatrix} = \begin{pmatrix} \sqrt{T} & -\sqrt{R} \\ \sqrt{R} & \sqrt{T} \end{pmatrix} \begin{pmatrix} C(l_+, t) \\ A(t) \end{pmatrix}. \quad (6.4)$$

The input field $A(t)$ satisfies the usual commutation relation,

$$\left[A(t), A^\dagger(t') \right] = \delta(t - t'). \quad (6.5)$$

6.1.2 Linearised response in the space-frequency domain

Fourier-transforming equations (6.1) - (6.3) from the time to the frequency domain ($t \rightarrow \omega$) and assuming the empty-cavity case, gives the following set of partial differential equations in the frequency domain:

$$\begin{aligned} -i\omega c(z, \omega) + v_g \partial_z c(z, \omega) &= 0 \\ c(0_+, \omega) &= \sqrt{R}c(l_-, \omega) + \sqrt{T}a(\omega) \\ b(\omega) &= \sqrt{T}c(l_-, \omega) - \sqrt{R}a(\omega). \end{aligned} \quad (6.6)$$

Here, $c(z, \omega)$, $a(\omega)$ and $b(\omega)$ denote the Fourier transforms of $C(z, t)$, $A(t)$ and $B(t)$, respectively. The first of these equations can be solved for the cavity field after one round trip:

$$c(l_-, \omega) = c(0_+, \omega)e^{i\omega\tau} \quad (6.7)$$

where τ is the cavity round-trip time. Combining equations (6.6) and (6.7) yields the cavity field in the frequency domain, i.e.

$$c(0_+, \omega) = \frac{\sqrt{T}}{1 - \sqrt{R}\exp(i\omega\tau)} a(\omega) = G(\omega)a(\omega), \quad (6.8)$$

where $G(\omega)$ is a Green's function which describes the linear response of the cavity. The

density of states inside the cavity is then given by

$$|G(\omega)|^2 = \frac{T}{(1 - \sqrt{R})^2 + 4\sqrt{R}\sin^2\left(\frac{\omega\tau}{2}\right)} \quad (6.9)$$

which is in fact the well-known transmission function of a Fabry-Perot-type, lossless cavity. This result demonstrates a fundamental aspect of cavity physics; the boundary conditions imposed by the cavity are responsible for a re-distribution of the density of states, leading to an enhancement of $|G(\omega)|^2$ near the cavity resonances^[295], as plotted in Figure 6.2. The total number of states is nevertheless conserved which can be easily verified by integrating $|G(\omega)|^2$ over multiples of the free spectral range. This result is at the very heart of the cavity approach for improving the signal-to-noise ratio in the Raman memory.




Figure 6.2: Modulus-square of the Green's function $G(\omega)$ for $\sqrt{R} = 0.75$ and $T = 1$.

Using equation (6.6), the output field $b(\omega)$ can also be expressed in terms of a Green's function G'_ω ,

$$b(\omega) = \frac{1 - \sqrt{R}\exp(-i\omega\tau)}{1 - \sqrt{R}\exp(i\omega\tau)} e^{i\omega\tau} a(\omega) = G'_\omega a(\omega), \quad (6.10)$$

Having derived an expression for $|G(\omega)|^2$, it is important to clarify the effect of a cavity-imposed re-distribution of the density of states on the memory and noise processes; the

reduction of the density of states at the anti-Stokes frequency is accompanied by a corresponding enhancement at the Stokes frequency. Since the signal field, which we wish to store, experiences this very same enhancement, the overall effect of the re-distribution is the suppression of noise *relative* to the signal, or equivalently, the enhancement of the signal above the noise floor. When referring to noise suppression, we will therefore always imply a comparison with the signal, since the metric of interest is ultimately the signal-to-noise ratio.

6.2 Cavity Raman memory equations

The following derivation of the dynamical equations governing the cavity Raman memory is based on work by Joshua Nunn. The Maxwell-Bloch equations describing the coherent Raman interaction of the Stokes field S and anti-Stokes field A with the spin wave coherence B were presented in Section 5.4.1. For convenience, these equations are repeated below, where we have made the coordinate transformation $\tau \rightarrow t$:

$$\begin{aligned} \left(\partial_z + \frac{1}{c}\partial_t\right) S(z, t) &= ik_S\Omega(z, t)B(z, t) + i\delta k S(z, t) \\ \left(\partial_z + \frac{1}{c}\partial_t\right) A(z, t) &= ik_A\Omega(z, t)B^\dagger(z, t) \\ \partial_t B(z, t) &= ik_S\Omega^*(z, t)S(z, t) + ik_A\Omega(z, t)A^\dagger(z, t) + iq|\Omega(z, t)|^2B(z, t). \end{aligned} \tag{6.11}$$

Here we have re-defined the coupling constants k_S and k_A for the Stokes and anti-Stokes fields to make explicit their dependence on the control Rabi frequency $\Omega(z, t)$. As before, δk is the phase mismatch of the FWM process and q parametrises the A.C. Stark shift induced by the control field.

As in the previous section, we will consider a ring cavity geometry, as shown in Figure 6.1,

with intensity reflection and transmission coefficients R and $T = 1 - R$. The boundary conditions for the Stokes and anti-Stokes fields at the coupling junction are given by

$$\begin{aligned} S(0, t) &= \sqrt{R}\eta e^{i\phi_S} S(l, t) - \sqrt{T}S_{\text{in}}(t) \\ A(0, t) &= \sqrt{R}\eta e^{i\phi_A} A(l, t) - \sqrt{T}A_{\text{in}}(t), \end{aligned} \tag{6.12}$$

where l is the round-trip length and $\phi_{S,A}$ is the phase accumulated by the fields as they propagate once around the cavity. We have also included intracavity loss by defining η as the amplitude transmission for each round-trip.

The Stokes field $S(l, t)$ after one pass of the cavity can be approximated using a Taylor expansion of $S(0, t)$ and substituting equation (6.11) for $\partial_z S(z, t)$,

$$\begin{aligned} S(l, t) &\approx S(0, t) + l\partial_z S(z, t) \Big|_{z=0} \\ &= S(0, t) - \frac{l}{c}\partial_t S(0, t) + ilk_S\Omega(0, t)B(0, t) + il\delta k S(0, t). \end{aligned} \tag{6.13}$$

This approximation assumes that the round-trip time $\tau = l/c$ is short compared to the pulse duration, or equivalently, that the spectral width of the cavity resonances is large compared to the pulse bandwidth. Intuitively, this can be understood as the requirement that the full spectrum of the incident signal should fit into the cavity's transmission window. As a result of this approximation, S varies slowly over the length of the cavity and we further assume that the round-trip phase mismatch is small, $\delta kl \ll 1$.

Likewise, the Taylor expansion of $A(l, t)$ gives

$$\begin{aligned} A(l, t) &\approx A(0, t) + l\partial_z A(z, t) \Big|_{z=0} \\ &= A(0, t) - \frac{l}{c}\partial_t A(0, t) + ilk_A\Omega(0, t)B^\dagger(0, t). \end{aligned} \tag{6.14}$$

Substituting the expressions for $S(l, t)$ and $A(l, t)$ into the boundary condition (6.12) gives the following set of partial differential equations for the intracavity fields $s(t) = \sqrt{\tau}S(0, t)$, $a(t) = \sqrt{\tau}A(0, t)$ and $b(t) = \sqrt{l}B(0, t)$:

$$\begin{aligned}\partial_t s &= -\gamma_- s + ic\delta k s + i\sqrt{c}k_S\Omega b - \kappa S_{\text{in}}, \\ \partial_t a &= \gamma_+ a + i\sqrt{c}k_A\Omega b + \kappa A_{\text{in}}, \\ \partial_t b &= i\sqrt{c}k_S\Omega^* s + i\sqrt{c}k_A\Omega a^\dagger + iq|\Omega|^2 b,\end{aligned}\tag{6.15}$$

where $\gamma_\pm = (1 \pm \eta\sqrt{R})/\eta\sqrt{R}\tau$ and $\kappa = \sqrt{T}/\sqrt{R}\eta\tau$. As the cavity is tuned to be resonant at the Stokes frequency and anti-resonant at the anti-Stokes frequency, the round-trip phases of the two fields have been set to $\phi_S = 0$ and $\phi_A = \pi$.

In the adiabatic limit of slowly-varying amplitudes, one can make the approximation $\partial_t s \approx \partial_t a \approx 0$. Solving for s and a gives

$$\begin{aligned}s &= \frac{i\sqrt{c}k_S\Omega}{\gamma_- - ic\delta k} b - \frac{\kappa}{\gamma_- - ic\delta k} S_{\text{in}}, \\ a &= -\frac{i\sqrt{c}k_A\Omega}{\gamma_+} b^\dagger - \frac{\kappa}{\gamma_+} A_{\text{in}}, \\ \partial_t b &= \left[iq - c \left(\frac{|k_S|^2}{\gamma_- - ic\delta k} + \frac{|k_A|^2}{\gamma_+} \right) \right] |\Omega|^2 b - \frac{i\sqrt{c}k_S\Omega^*\kappa}{\gamma_- - ic\delta k} S_{\text{in}} - \frac{i\sqrt{c}k_A\Omega\kappa}{\gamma_+} A_{\text{in}}^\dagger.\end{aligned}\tag{6.16}$$

As in Section 2.2, we will proceed by transforming to the quasi-time coordinate $\omega(t) = \int_{-\infty}^t |\Omega(t')|^2 dt'$. The partial derivative becomes $\partial_t = |\Omega(t)|^2 \partial_\omega$. Equations (6.16) are then

given by

$$\begin{aligned}
\sigma &= \frac{i\sqrt{c}k_S}{\gamma_- - ic\delta k} b - \frac{\kappa}{\gamma_- - ic\delta k} \sigma_{\text{in}}, \\
\alpha &= -\frac{i\sqrt{c}k_A}{\gamma_+} b^\dagger - \frac{\kappa}{\gamma_+} \alpha_{\text{in}}, \\
\partial_\omega b &= \left[iq - c \left(\frac{|k_S|^2}{\gamma_- - ic\delta k} + \frac{|k_A|^2}{\gamma_+} \right) \right] b - \frac{i\sqrt{c}k_S^* \kappa}{\gamma_- - ic\delta k} \sigma_{\text{in}} - \frac{i\sqrt{c}k_A \kappa}{\gamma_+} \alpha_{\text{in}}^\dagger,
\end{aligned} \tag{6.17}$$

where the Stokes and anti-Stokes fields have been normalised by the Rabi frequency, i.e.

$$\sigma = s/\Omega \text{ and } \alpha = a/\Omega.$$

Solving for the spin wave $b(\omega)$ gives

$$b(\omega) = b(0)e^{f\omega} - g_\sigma \int_0^\omega e^{f(\omega-\omega')} \sigma_{\text{in}}(\omega') d\omega' - g_\alpha \int_0^\omega e^{f(\omega-\omega')} \alpha_{\text{in}}^\dagger(\omega') d\omega', \tag{6.18}$$

where

$$f = \left[iq - c \left(\frac{|k_S|^2}{\gamma_- - ic\delta k} + \frac{|k_A|^2}{\gamma_+} \right) \right], \tag{6.19}$$

and

$$g_\sigma = \frac{i\sqrt{c}k_S^* \kappa}{\gamma_- - ic\delta k} \quad g_\alpha = \frac{i\sqrt{c}k_A \kappa}{\gamma_+}. \tag{6.20}$$

Using equations (6.17) and (6.18), we can now determine the output fields emerging from the cavity. Experimentally, it is convenient to measure the fields transmitted through the cavity, rather than those reflected at the input junction. The reason for this lies in the fact that the output intensity of a cavity operated in transmission is a direct indicator of how well the cavity is aligned, i.e. maxima in transmission are more readily identified than minima in reflection. While experimentally simpler, this configuration does, however, introduce a source of loss, as the field retrieved from the memory can leak through the input coupler of the cavity. Although cavity operation in transmission may limit the attainable memory

efficiency, any transmitted FWM noise will be equally affected so that noise suppression is not expected to deteriorate. With this in mind, we will now consider the Stokes field S_{trans} transmitted through the output coupler of a two-junction cavity, as shown in Figure 6.4. Having determined the intracavity field $S(l', t)$ at the position of the output coupler, S_{trans} is simply given by

$$S_{\text{trans}} = \sqrt{T_2} S(l', t) = \sqrt{T_2} \eta' e^{i\phi'_S} S(0, t) \approx \sqrt{T_2} S(0, t), \quad (6.21)$$

where T_2 is the intensity transmission coefficient of the output coupler and l' is the optical beam path between the input and output couplers. In the last step we have assumed that the partial intracavity loss $1 - \eta' \approx 0$ and have set the phase $\phi'_S = 0$.

Likewise, for the transmitted anti-Stokes field,

$$A_{\text{trans}} = \sqrt{T_2} A(l', t) = \sqrt{T_2} \eta' e^{i\phi'_A} A(0, t) \approx \sqrt{T_2} A(0, t). \quad (6.22)$$

Transforming to the quasi-time coordinate $\omega(t)$ and inserting $b(\omega)$ into the expressions for the intracavity modes σ and α , we find

$$\begin{aligned} \sigma_{\text{trans}} &= \sqrt{T_2} \left[p_\sigma e^{f\omega} b(0) - p_\sigma g_\sigma \int_0^\omega e^{f(\omega-\omega')} \sigma_{\text{in}}(\omega') d\omega' - p_\sigma g_\alpha \int_0^\omega e^{f(\omega-\omega')} \alpha_{\text{in}}^\dagger(\omega') d\omega' - v_\sigma \sigma_{\text{in}} \right] \\ \alpha_{\text{trans}} &= \sqrt{T_2} \left[p_\alpha e^{f^*\omega} b(0) - p_\alpha g_\sigma^* \int_0^\omega e^{f^*(\omega-\omega')} \sigma_{\text{in}}^\dagger(\omega') d\omega' - p_\alpha g_\alpha^* \int_0^\omega e^{f^*(\omega-\omega')} \alpha_{\text{in}}(\omega') d\omega' - v_\alpha \alpha_{\text{in}} \right], \end{aligned} \quad (6.23)$$

where we have defined $p_\sigma = \frac{i\sqrt{ck_S}}{\gamma_- - ic\delta k}$, $v_\sigma = \frac{\kappa}{\gamma_- - ic\delta k}$, $p_\alpha = -\frac{i\sqrt{ck_A}\Omega}{\gamma_+}$ and $v_\alpha = \frac{\kappa}{\gamma_+}$.

Note that the reflected output fields can be derived in a similar manner using the bound-

ary conditions at the input coupler which are given by

$$\begin{aligned} S_{\text{ref}}(t) &= \sqrt{R}S_{\text{in}}(t) + \sqrt{T}\eta e^{i\phi_S} S(l, t) \\ A_{\text{ref}}(t) &= \sqrt{R}A_{\text{in}}(t) + \sqrt{T}\eta e^{i\phi_A} A(l, t). \end{aligned} \quad (6.24)$$

Finally, the retrieved field can be found by replacing the initial spin wave $b(0)$ with the final spin wave $b(W)$ at the end of the storage process, where $W = \int_{-\infty}^{\infty} |\Omega(t')|^2 dt'$:

$$\begin{aligned} \sigma_{\text{retr}}(\omega) &= -\sqrt{T_2} p_\sigma e^{f\omega} \left[g_\sigma \int_0^W e^{f(W-\omega')} \sigma_{\text{in}}(\omega') d\omega' + g_\alpha \int_0^W e^{f(W-\omega')} \alpha_{\text{in},1}^\dagger(\omega') d\omega' \right] \\ &\quad - \sqrt{T_2} p_\sigma g_\alpha \int_0^\omega e^{f(\omega-\omega')} \alpha_{\text{in},2}^\dagger(\omega') d\omega'. \end{aligned} \quad (6.25)$$

Here, any fields that are initially in the vacuum state, i.e. $\sigma_{\text{in},2}$, have been dropped. We will refer back to this equation when examining the modes coupled by the linear mapping of the cavity Raman process in Section 6.4.

6.3 Cavity enhancement

Compared to free-space implementations, memory operation inside a cavity benefits from an increased effective optical depth due to enhanced coherent coupling between the signal field and the atomic ensemble^{[58] [296] [297]}. This effective optical depth can be quantified by the cooperativity parameter D which accounts for the cavity-induced enhancement of the coupling strength. The optimal memory efficiency is uniquely determined by the value of D , independent of the detuning Δ and the temporal shape of the control pulse $\Omega(t)$. We have already encountered this very same result in our previous discussion of the free-space model in Section 2.2.1 for which the equivalent figure of merit is the optical depth d . This

result applies to adiabatic memory schemes – both far-off-resonant Raman and resonant EIT protocols – as well as non-adiabatic photon-echo memories.

As alluded to in Section 2.2.1, the sole dependence of the efficiency on D (or d in free space) reflects the fact that the branching ratio between collectively enhanced emission into the optical output mode and spontaneous decay into undesired modes is determined by the optical depth alone. This was developed by Alexey Gorshkov^[298], alongside the pivotal observation that optimal storage is the exact time-reversal of optimal retrieval^{[117][186]} and $\eta_{\text{stor,opt}} = \eta_{\text{ret,opt}}$. While a formal proof of this result is given in reference^[208], it can be intuitively understood as a manifestation of the \mathbb{T} -symmetry of the unitary beam-splitter relations (see Section 2.2.4).

The input-output relation for the signal field $S(t)$ can be re-cast in terms of the cavity decay rate $\Gamma_{\text{cav}} = 2\pi\delta\nu$, where $\delta\nu$ is the resonance bandwidth. As shown in reference^[171], the equations of motion for the cavity memory interaction in the adiabatic approximation are dependent on the cooperativity parameter^[299] $D = |\kappa|^2/(\Gamma_{\text{cav}}\gamma)$, equivalent to d in equation (2.38). As before, κ is the coupling constant between the signal mode and the atomic polarisation, and γ is the spontaneous emission rate from the excited state.

In order to make explicit the cavity-dependent modification of the optical depth, we can re-write D in terms of the free-space optical depth $d = |\kappa|^2l/(c\gamma)$, as

$$D = d \times \left(\frac{1}{\Gamma_{\text{cav}}} \times \frac{c}{l} \right) = d \times \left(\frac{1}{2\pi\delta\nu} \times 2\Delta\nu \right) = d \times \left(\frac{\Delta\nu}{\delta\nu} \frac{1}{\pi} \right) = d \times \left(\frac{F}{\pi} \right). \quad (6.26)$$

where l is the ensemble length and we have defined the cavity free spectral range as $\Delta\nu = c/(2l)$ (see Section 6.7). Note that for simplicity we have assumed that the cavity length is equal to l . Equation (6.26) shows that the optical depth is enhanced by a factor proportional

to the cavity finesse $F = \Delta\nu/\delta\nu$. This factor, F/π , is in fact equal to the ratio of the cavity lifetime $1/(\Gamma_{\text{cav}}) = 1/(2\pi\delta\nu)$ to half the round-trip time $l/c = 1/(2\Delta\nu)$. The optical depth inside the cavity is therefore increased by the number of passes that a photon would make through the interaction region before leaking out^[298].

The retrieval efficiency can be computed in analogy with equation (2.42) in Chapter 2, as the ratio between the number of retrieved photons and the number of stored excitations, i.e.

$$\eta_{\text{ret}} = \frac{\int_{-\infty}^{\infty} |S_{\text{out}}(\tau)|^2 d\tau}{\int_{-\infty}^{\infty} |B_{\text{mem}}(\tau)|^2 d\tau}. \quad (6.27)$$

Provided that spin wave decoherence can be ignored and that no excitations remain at $t = \infty$ ($B_{\text{mem}}(\infty) = 0$), the optimal retrieval efficiency reduces to $D/(1 + D)$, independent of the control field. This necessarily assumes that the control is sufficiently strong to convert the stored spin wave fully into an optical output mode. According to the time-reversal argument, the maximum total memory efficiency for storage followed by retrieval is therefore given by $D^2/(1 + D)^2$.

6.4 Mode selectivity

In discussing the Raman interaction in Section 2.2.1, we introduced the linear optical mapping which connects the input signal S_{in} with the retrieved signal S_{out} :

$$S_{\text{out}}(t) = \int K(t, t') S_{\text{in}}(t') dt'. \quad (6.28)$$

The mapping is represented by the Green's function K which couples the j^{th} input mode ϕ_j to the j^{th} output mode ψ_j , where the coupling strength is given by the corresponding

singular value λ_j and all modes have been normalised.

$$K(t, t') = \sum_j \lambda_j \psi_j(t) \phi_j^*(t') = \lambda_1 \psi_1(t) \phi_1^*(t') + \lambda_2 \psi_2(t) \phi_2^*(t') + \dots \quad (6.29)$$

By convention, the terms in the above sum are ordered according to the magnitude of the singular values, with λ_1 being the largest.

If $\lambda_j = 0$ for $j \neq 1$, the memory interaction can only couple a single mode and $K(t, t') = \lambda_1 \psi_1(t) \phi_1^*(t')$. In this case, the linear transformation $K(t, t')$ is a separable function in t and t' and we can write

$$S_{\text{out}}(t) = \xi \times \psi_1(t), \quad (6.30)$$

where

$$\xi = \lambda_1 \int \phi_1^*(t') S_{\text{in}}(t') dt' \quad (6.31)$$

describes the overlap between the signal S_{in} and the input mode function ϕ_1 of the linear map $K(t, t')$. The efficiency with which the input field is mapped to the output field therefore depends on the value of ξ . If S_{in} is orthogonal to ϕ_1 , then the integral in equation (6.31) is zero and $\eta = 0$. Consequently, a single mode memory is maximally mode selective, since any incident signal mode orthogonal to ϕ_1 remains unchanged by the interaction. Mode selectivity is a useful functionality, and indeed a necessary prerequisite for applications aiming to deploy temporal-mode multiplexing^{[300][301]}.

Generally, as the efficiency of the memory is increased, additional modes are coupled by the interaction due to the increasing magnitude of the corresponding singular values^[121]. These terms can no longer be ignored in the mode expansion of $K(t, t')$ which cannot therefore be treated as a separable function. The resulting trade-off between coupling efficiency

and mode selectivity is discussed in references^[302]^[303].

The Raman memory inside a travelling-wave cavity can indeed be considered as a generalisation of this approach. The coupling strength of each individual round-trip is comparatively weak so that $\lambda_1 \gg \lambda_{j \neq 1}$. However, multiple passes of the ring cavity interfere constructively to give high overall efficiency at the output, provided that the signal is resonant with the cavity mode. As a result, the memory interaction is simultaneously efficient and highly mode selective. This particularly interesting feature of the cavity Raman memory was first examined by Josh Nunn.

The linear transformation relating the input signal field and the retrieved field (equation (6.25)) contains an explicit expression for the Green's function K . Ignoring all noise contributions, the mapping can be expressed, in terms of the quasi-time coordinate $\omega(t)$, as

$$\sigma_{\text{out}}(\omega) = \int K(\omega, \omega') \sigma_{\text{in}}(\omega') d\omega' \quad (6.32)$$

where

$$K(\omega, \omega') = -\sqrt{T_2} g_{\sigma} p_{\sigma} e^{f\omega} e^{f(W-\omega')}. \quad (6.33)$$

$K(t, t')$ can be recovered via the coordinate transformation $(\omega, \omega') \rightarrow (t, t')$, showing that $K(t, t')$ is separable in t and t' . Hence, the memory interaction is exactly single-mode, provided that the approximations used to derive equation (6.25) are valid. Apart from noise suppression, the cavity Raman memory therefore also provides high mode selectivity independent of efficiency.

6.5 Noise suppression

The coherent model of four-wave mixing noise, introduced in the previous chapter, can be used to examine the effect of noise suppression on the photon statistics of the retrieved field following storage of single photons in the Raman memory. For this purpose, we will compare the coupling strength of the memory interaction, $C = \sqrt{Wd\gamma}/\Delta$, with that of the FWM process seeded by anti-Stokes scattering, given by $C' = \sqrt{Wd\gamma}/(\Delta + \Delta_{\text{HFS}})$. Figure 6.3 shows the predicted $g^{(2)}$ autocorrelation (equation 5.13), as the ratio of coupling strengths $R = C'/C$ is artificially varied. In free space, R is fixed by the ratio of detunings of the anti-Stokes and Stokes fields, i.e.

$$R = \frac{C'}{C} = \frac{\Delta}{\Delta + \Delta_{\text{HFS}}} \approx \frac{15}{15 + 9} = 0.625. \quad (6.34)$$




Figure 6.3: Theoretical prediction for the $g^{(2)}$ -autocorrelation of the retrieved field following storage of a heralded SPDC input photon, as the ratio of coupling strengths $R = C'/C$ is artificially varied. The vertical dotted line indicates $R = 0.625$, corresponding to the ratio of detunings of the anti-Stokes and Stokes fields.

As described in Section 6.1.2, the cavity boundary conditions impose a re-distribution of the density of states, with an enhancement near the cavity resonances and a corresponding

reduction between resonance peaks. If the cavity is tuned to be off-resonant at the anti-Stokes frequency, emission at this frequency is strongly suppressed due to the reduction in the density of states g_{AS} .

For memory operation inside a cavity, the ratio R can therefore be re-defined as

$$R = \frac{C'}{C} = \frac{g_{AS}}{g_S} \frac{\Delta}{\Delta + \Delta_{HFS}}, \quad (6.35)$$

where g_S is the density of states at the signal (Stokes) frequency.

Suppressing the anti-Stokes coupling by a factor of 2.5 would result in a reduction of R sufficient to observe a $g^{(2)} < 1$ and preserve non-classical statistics. With a suppression factor of 30, the statistics of the output field would be as non-classical as those of the input signal¹.

6.6 The ring cavity design

In the previous section, we derived the transfer function of a bare ring cavity by considering the beam splitter equations relating the input field with the intracavity modes and Fourier-transforming these equations from the time domain to the frequency domain. The cavity's transmission spectrum was then determined from the square modulus of the Green's function which is equivalent to the density of states inside the cavity.

In the following section I will present an alternative approach to deriving the transmission function. Before doing so, I would like to motivate the choice of cavity geometry which has so far been underlying our discussion.

Implementing a Raman memory inside a linear Fabry-Perot-type cavity has a few consid-

¹Noise originating from the SPDC source would outweigh any noise contribution from the memory.

erable disadvantages. Firstly, the standing wave established inside the linear cavity results in intensity nodes along the cavity axis. While the total intracavity intensity is the same as for other cavity geometries, the node structure would severely reduce the memory lifetime due to longitudinal atomic diffusion out of the interaction region. The maximum storage time would therefore be limited by the transit time of an atom across one intensity anti-node.

Secondly, a linear configuration – by definition – does not allow unidirectional propagation. Hence, it would not be possible to introduce the pump laser into the cavity in a counter-propagating direction relative to the signal and control pulses, as has been the strategy with the free-space memory. These problems can be circumvented with a travelling-wave ring cavity, the geometry of which is shown in Figure 6.7.

6.7 The lossy ring cavity

The field amplitude at the output ports of a cavity can be determined classically by summing the contributions from each round-trip, taking into account both the accumulation of phase^[304] as well as the reduction in amplitude due to loss. The phase associated with one complete pass of a cavity of round-trip length $2l$ is given by

$$\delta = kx = \frac{2\pi}{\lambda} \times 2nl = \frac{4\pi ln}{\lambda} = \frac{4\pi n l \nu}{c}, \quad (6.36)$$

where the frequency $\nu = c/\lambda$ and n is the refractive index of the intracavity medium. For the ring cavity shown in Figure 6.4, the field transmitted through the output coupler is given by the sum

$$E_{\text{trans}} = (t_1 t_2 + t_1 r_2 r_1 t_2 \eta e^{i\delta} + t_1 r_2^2 r_1^2 t_2 \eta^2 e^{2i\delta} + \dots) \times E_{\text{in}} \quad (6.37)$$

where $r_{1/2}$ is the amplitude reflection coefficient of the input/output coupler² and $t_{1/2}$ is the corresponding amplitude transmission coefficient which satisfies the relation $|r_{1/2}|^2 + |t_{1/2}|^2 = 1$. We have also defined the loss parameter η such that the fractional intracavity intensity loss for one round-trip is given by $1 - |\eta|^2$. Note that, as before, we have ignored any partial loss $1 - |\eta'|^2$ between the input and output couplers which are separated by $l' \ll 2l$ and set the phase $\delta' = \frac{2\pi}{\lambda} \times l' = 0$.




Figure 6.4: Two-junction ring cavity (total length $2l$) with amplitude reflection and transmission coefficients $r_{1/2}$ and $t_{1/2}$, respectively. The overall intracavity loss is given by $1 - |\eta|^2$. The separation l' between the input and output couplers is taken to be sufficiently small so that any partial loss can be ignored over this distance.

Hence, the amplitude transmission function, normalised by the incident field amplitude, is given by

$$\frac{E_{\text{trans}}}{E_{\text{in}}} = \sum_{m=0}^{\infty} (r_1 r_2)^m \eta^m e^{i\delta m} t_1 t_2 = \frac{t_1 t_2}{1 - r_1 r_2 \eta e^{i\delta}}. \quad (6.38)$$

The form of the corresponding intensity transmission function evokes the familiar Airy function of a planar Fabry-Perot cavity, with $\zeta = r_1 r_2 \eta$ replacing the reflectance R ,

$$\left| \frac{I_{\text{trans}}}{I_{\text{in}}} \right| = \frac{|t_1 t_2|^2}{1 - 2r_1 r_2 \eta \cos \delta + (r_1 r_2 \eta)^2} = \frac{|t_1 t_2|^2}{(1 - \zeta)^2 + 4\zeta \sin^2 \left(\frac{\delta}{2} \right)}. \quad (6.39)$$

Up to a factor of $|t_2|^2$, equation (6.39) is identical to the intracavity density of states

²We will use lower-case letters to represent amplitude coefficients.

$|G(\omega)|^2$, as derived in Section 6.1.2. The resonance frequencies of the cavity are found by setting $\delta/2 = \pi m$, giving $\nu_{\text{res}} = m \times \frac{c}{2nl}$, where m is an integer. The free spectral range (FSR), defined as the frequency separation $\Delta\nu$ between successive resonance peaks, is inversely proportional to the cavity length, since $\Delta\nu = \frac{c}{2nl}$. The intensity of the intracavity field can be obtained by dividing equation (6.39) by the intensity transmission coefficient of the output port $T_2 = |t_2|^2$.

We have already encountered the cavity finesse in Section 6.3, defined as the ratio between the FSR and the FWHM bandwidth $\delta\nu$ of the transmission peaks,

$$F = \frac{\Delta\nu}{\delta\nu} = \frac{\pi}{2 \sin^{-1} \left(\frac{1-\zeta}{2\sqrt{\zeta}} \right)}. \quad (6.40)$$

Equation (6.40) shows that the finesse is fully determined by the reflectivity of the cavity mirrors and the intracavity loss, independent of the cavity length. As we shall see in Section 6.10.1, the parameters chosen for the ring cavity place it in the low-finesse regime.

We can similarly find an expression for the field reflected at the input mirror, i.e.

$$\frac{E_{\text{refl}}}{E_{\text{in}}} = r_1 - (t_1^2 r_2 \eta e^{i\delta} + t_1^2 r_2^2 r_1 \eta^2 e^{2i\delta} + \dots) = r_1 - t_1^2 r_2 e^{i\delta} \eta \sum_{m=0}^{\infty} r_1^m r_2^m \eta^m e^{im\delta} = \frac{r_1 - r_2 \eta e^{i\delta}}{1 - r_1 r_2 \eta e^{i\delta}}. \quad (6.41)$$

The reflected intensity is then given by the square modulus of equation (6.41),

$$\left| \frac{I_{\text{refl}}}{I_{\text{in}}} \right| = \frac{|r_1|^2 + |r_2 \eta|^2 - 2r_1 r_2 \eta \cos \delta}{1 + |r_1 r_2 \eta|^2 - 2r_1 r_2 \eta \cos \delta}. \quad (6.42)$$

On resonance, the transmitted and reflected intensities reduce to

$$\left| \frac{I_{\text{trans}}}{I_{\text{in}}} \right|_{\text{res}} = \frac{(1 - |r_1|^2)(1 - |r_2|^2)}{(1 - r_1 r_2 \eta)^2} \quad \left| \frac{I_{\text{refl}}}{I_{\text{in}}} \right|_{\text{res}} = \left| \frac{r_1 - r_2 \eta}{1 - r_1 r_2 \eta} \right|^2. \quad (6.43)$$

In the case of a lossless cavity with $\eta = 1$, it is straightforward to verify that energy is conserved, i.e. $\left| \frac{I_{\text{trans}}}{I_{\text{in}}} \right|_{\text{res}} + \left| \frac{I_{\text{refl}}}{I_{\text{in}}} \right|_{\text{res}} = 1$. The introduction of loss through $\eta < 1$ is illustrated in Figure 6.5 which shows the transmitted and reflected intensity spectra for the choice of cavity parameters discussed in Section 6.10.1.




Figure 6.5: Intensity spectra of the transmitted (blue) and reflected fields (red) as a function of frequency. The sum of these spectra is shown in green.

The condition for critical cavity coupling is satisfied when $r_1 - r_2\eta = 0$ such that the reflected modes interfere destructively at the input port, yielding $\left| \frac{I_{\text{refl}}}{I_{\text{in}}} \right|_{\text{res}} = 0$. Requiring critical coupling is therefore equivalent to maximising the intracavity field intensity – a key factor in enabling memory operation inside the cavity.

6.8 Extinction

One of the principal figures of merit guiding our design considerations is the amount of cavity-induced noise extinction. For this purpose, we will define a cavity extinction ratio Υ which is a useful parameter for assessing the suppression of the anti-Stokes frequency with respect to the resonant enhancement of the signal/Stokes frequency. As noted earlier, the overall result is a corresponding increase in the signal-to-noise ratio.

An upper bound on Υ is given by the ratio between the minima and maxima of the cavity

transmission function, corresponding to the optimal amount of noise suppression possible for CW fields. In this limit, the cavity extinction ratio is equal to

$$\Upsilon(\zeta) = \frac{|I_{\text{trans}}|_{\text{min}}}{|I_{\text{trans}}|_{\text{max}}} = \frac{(1 - \zeta)^2}{(1 + \zeta)^2}. \quad (6.44)$$

Equation (6.44) shows that $\Upsilon = \Upsilon(\zeta)$ only depends on $\zeta = r_1 r_2 \eta$, i.e. the reflectivities of the cavity mirrors and the intracavity loss. In other words, the cavity finesse $F(\zeta)$ ultimately limits the potential noise suppression. Since the finesse is inversely proportional to the cavity bandwidth, the requirements for broadband memory operation and large signal-to-noise ratio are in direct competition with each other. Note that this could be mitigated by increasing the free-spectral range of the cavity – a discussion we will return to in Section 6.10.1

We can also express $\Upsilon(\zeta)$ in terms of the cavity visibility $V(\zeta)$ as $\Upsilon(\zeta) = (1 - V(\zeta))/(1 + V(\zeta))$, where

$$V(\zeta) = \frac{|I_{\text{trans}}|_{\text{max}} - |I_{\text{trans}}|_{\text{min}}}{|I_{\text{trans}}|_{\text{max}} + |I_{\text{trans}}|_{\text{min}}} = \frac{(1 + \zeta)^2 - (1 - \zeta)^2}{(1 + \zeta)^2 + (1 - \zeta)^2}. \quad (6.45)$$

The finite width of the pulses can be taken into account by evaluating the overlap between the cavity transmission function and a Gaussian pulse with FWHM $\delta\nu$. For the case of a critically coupled cavity, we can re-write the transmission function in equation (6.39) in terms of ζ , i.e.

$$T(\nu, \zeta, l) = \frac{(1 - \zeta)(1 - \zeta/\eta^2)}{(1 - \zeta)^2 + 4\zeta \sin^2\left(\frac{2\pi l\nu}{c}\right)}. \quad (6.46)$$

The overlap $\Theta(\nu, \zeta, l)$ between $T(\nu, \zeta, l)$ and a Gaussian function $f(\nu, \nu_0, \delta\nu)$, centred at ν_0 , is then given by

$$\Theta(\zeta, l, \nu_0, \delta\nu) = \int_{-\infty}^{\infty} T(\nu, \zeta, l) \times f(\nu, \nu_0, \delta\nu) d\nu. \quad (6.47)$$

Using equation (6.47), we can define an effective cavity extinction ratio $\Upsilon'(\zeta, \delta\nu) = \Theta(\zeta, \nu_a, \delta\nu)/\Theta(\zeta, \nu_s, \delta\nu)$, where we have assumed that the cavity is resonant at the signal frequency ν_s and anti-resonant at the anti-Stokes frequency $\nu_a = \nu_s + 2\Delta_{\text{HFS}}$. Experimentally, $\Upsilon'(\zeta, \delta\nu)$ can be estimated from the measured visibility $V'(\zeta, \delta\nu)$ of the cavity transmission function which is defined as

$$V'(\zeta, \delta\nu) = \frac{\Theta(\zeta, \nu_s, \delta\nu) - \Theta(\zeta, \nu_a, \delta\nu)}{\Theta(\zeta, \nu_s, \delta\nu) + \Theta(\zeta, \nu_a, \delta\nu)}. \quad (6.48)$$

Note that the measured visibility is the result of a convolution between the cavity response function and the incident pulse shape and is therefore dependent on the pulse bandwidth $\delta\nu$. After re-arranging equation (6.48), we obtain

$$\Upsilon'(\zeta, \delta\nu) = \frac{1 - V'(\zeta, \delta\nu)}{1 + V'(\zeta, \delta\nu)}. \quad (6.49)$$

The discrepancy between the ideal and effective ratios Υ and Υ' is shown in Figure 6.6 as a function of the reflectivity R_2 of the output coupler. In evaluating the cavity extinction ratios we have assumed a Gaussian function with $\delta\nu = 1.2$ GHz and anticipated the choice of cavity parameters as well as the measured loss parameter η . We will discuss this choice of parameters in Section 6.10, but before doing so, we will first examine the geometric requirements for a stable cavity.




Figure 6.6: Ideal and effective cavity extinction ratios Υ (red) and Υ' (blue) - taking into account a Gaussian pulse bandwidth of 1.2 GHz - as a function of R_2 , while maintaining critical cavity coupling.

6.9 Ring cavity geometry

Figure 6.7 illustrates the isosceles geometry of the ring cavity which consists of an input coupling mirror, a concave mirror and an output coupling prism.




Figure 6.7: Triangular ring cavity (not to scale). The total round-trip path length is given by $d_1 + d_2$ and the separation between the input and output couplers is equal to $2x$. This notation will be used in Section 6.9.1.

A triangular shape was chosen to maximise the space available for the memory cell, while minimising the round-trip length of the cavity. The reasoning behind this will be discussed in Section 6.10.1. Given this triangular geometry, the curvature of the concave mirror had to be chosen such that the propagating fields overlap after each round-trip. This stability

criterion can be examined by tracing the incident beam, assumed to be Gaussian, around the cavity via the method of ray transfer matrix analysis.

6.9.1 Gaussian beam propagation

A prerequisite of a stable ring cavity is the condition that a propagating beam is returned to its original geometry after each round-trip. In order to ensure this, we will consider the set of Gaussian beam parameters, i.e. the beam radius $w(z)$ and the radius of curvature $R(z)$, where z is the distance from the beam waist along the beam axis. These parameters are given by

$$R(z) = z + \frac{z_R^2}{z} \quad \text{and} \quad w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \quad (6.50)$$

where w_0 is the beam waist. The Rayleigh range $z_R = \frac{\pi w_0^2}{\lambda}$ describes the distance from the waist over which the cross-sectional area of the beam doubles in size³. The beam radius and radius of curvature can be combined in the complex beam parameter $q(z)$ which fully characterises the beam and is given by

$$\frac{1}{q(z)} = \frac{1}{R(z)} - i \frac{\lambda}{\pi w^2(z)}. \quad (6.51)$$

As the Gaussian beam propagates through an optical system, the corresponding change in $q(z)$ can be calculated through first-order ray transfer matrix analysis which assumes the paraxial approximation of small angles relative to the beam axis. The ray transfer matrix $M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$ describes the optical system defined between an input plane and an output plane which are both perpendicular to the optical axis. The complex beam parameter q_f ,

³The confocal parameter is a related quantity, defined as twice the Rayleigh range.

after propagation through this system, is related to the initial q -parameter q_i through

$$q_f = \frac{A \cdot q_i + B}{C \cdot q_i + D} \quad (6.52)$$

Requiring $q_f = q_i$ yields a quadratic equation with solutions

$$p = -\frac{A - D}{2B} \pm \frac{\sqrt{\left(\frac{A+D}{2}\right)^2 - 1}}{B} \quad (6.53)$$

where we have re-defined the beam parameter as $p = 1/q$. As derived in Appendix B, the cavity is stable provided that

$$-1 < \frac{1}{2} \text{Tr}(\text{M}) < +1 \quad (6.54)$$

or alternatively,

$$0 < \frac{1}{4} \text{Tr}(\text{M}) + \frac{1}{2} < +1 \quad (6.55)$$

$$0 < g_1 g_2 < 1$$

where, in the second line, the condition has been re-expressed in terms of the stability parameters g_1 and g_2 . This notation is commonly used to describe the stability of cavities consisting of two curved mirrors. For this particular case, it is straightforward to show that $g_{1/2} = 1 - L/R_{1/2}$ where $R_{1/2}$ is the curvature of the first/second mirror and L is their separation. The stability criterion can be illustrated graphically by plotting g_1 against g_2 so that stable cavity configurations are represented by the areas bounded by the line $g_1 g_2 = 1$ and the coordinate axes (see Figure 6.11 below).

The Rayleigh range z_R of the Gaussian cavity mode was chosen to be equal to half the round-trip length of the ring cavity pictured in Figure 6.7, i.e. $z_R = l$, with the beam waist located at the mid-point between the input and output couplers. This choice ensures limited

divergence of the intracavity fields over the length of the memory interaction region.

The transfer matrix for propagation over a distance l is given by $M_1(l) = \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix}$. If the beam propagates through a series of optical materials with different refractive indices, then $l = \sum_i l_i/n_i$, where n_i is the refractive index associated with a length of material l_i and the sum runs over the number of different materials⁴. Reflection from a curved mirror and subsequent propagation is described by the matrix $M_2(l) = \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -2/R_c & 1 \end{pmatrix}$ where R_c is the radius of curvature.

In order to trace the change in beam parameters along the cavity path, it is convenient to *unfold* the ring cavity, as illustrated in Figure 6.8.

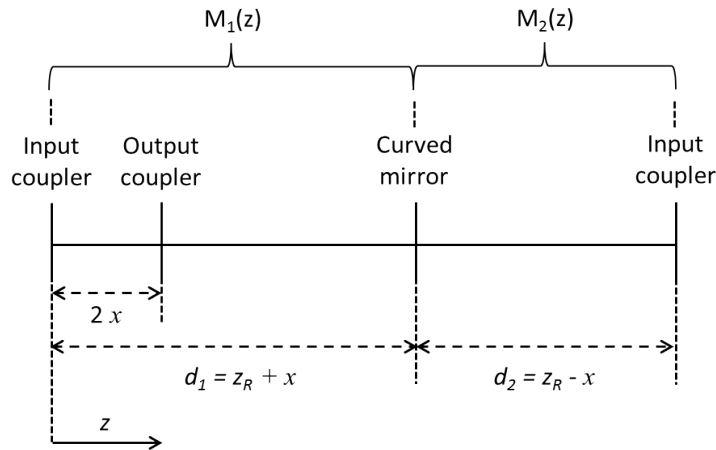


Figure 6.8: Unfolded ring cavity, showing the ray transfer matrices A_1 and A_2 for propagation from $z = 0$ to $z = d_1$ and from $z = d_1$ to $z = d_2$, respectively.

The path length between the input coupler and the curved mirror is defined as $d_1 = z_R + x$ where x is half the distance between the input and output mirrors. The remaining path length necessary to complete a round-trip is equal to $d_2 = z_R - x$ (see Figure 6.7). The

⁴This can be verified by multiplying the matrices for free-space propagation and refraction at a flat interface, and setting $n_i = n_{\text{air}} = 1$:

$$\begin{pmatrix} 1 & 0 \\ 0 & \frac{n_f}{n_i} \end{pmatrix} \begin{pmatrix} 1 & l_{\text{free}} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \frac{n_i}{n_f} \end{pmatrix} = \begin{pmatrix} 1 & \frac{l_{\text{free}} n_i}{n_f} \\ 0 & 1 \end{pmatrix} \quad (6.56)$$

ray transfer matrices corresponding to one round-trip are therefore given by $M_1(z)$ for $0 < z < d_1$ and $M_2(z)$ for $0 < z < d_2$. The beam size and radius of curvature are calculated from

$$w(z) = \frac{1}{\sqrt{\frac{-\pi \Im\left(\frac{1}{q(z)}\right)}{\lambda}}} \quad \text{and} \quad R(z) = \frac{1}{\Re\left(\frac{1}{q(z)}\right)}. \quad (6.57)$$

Figure 6.9 shows $w(z)$ and $R(z)$ as a function of z for one round trip, i.e. $0 < z < d_1 + d_2$:

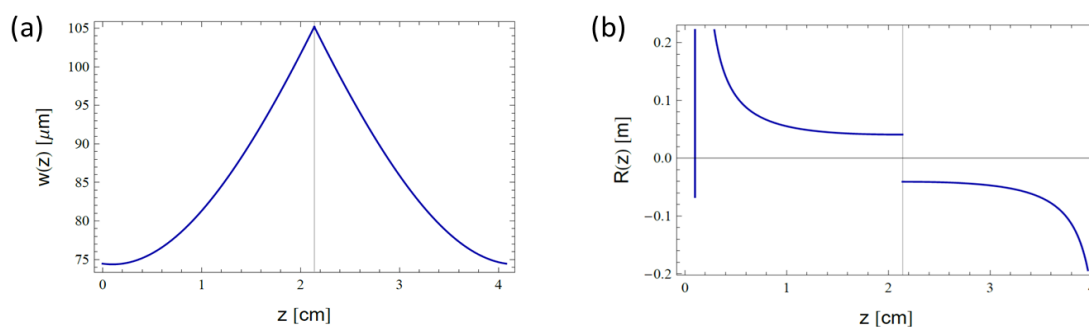


Figure 6.9: (a) Beam radius $w(z)$ for one round-trip. The vertical line indicates the position of the curved mirror at $z = d_1$. (b) Radius of curvature $R(z)$ for one round-trip. Notice that the sign changes upon reflection from the curved mirror at $z = d_1$.

From equation (6.57) we can determine the optimal beam parameters for stable propagation around the cavity. Given the choice of z_R , the radius of the incident beam at the input coupler should equal $w(0) = 76.6 \mu\text{m}$, with an optimal beam waist of $w_{\text{opt}} = w(x) = 74.4 \mu\text{m}$, located at $z = x$ between the input and output couplers. At the concave mirror ($z = d_1$) the beam expands to its largest radius, given by $w(d_1) = 105.2 \mu\text{m}$, before being re-focussed to $w(x)$ during the second half of the round-trip. Figure 6.10 shows $w(z)$ and $R(z)$ after one round-trip ($z = d_1 + d_2$), as a function of the radius R_c of the curved mirror. Clearly, the beam is only returned to its initial size after one round-trip if the mirror radius is equal to twice the Rayleigh range, $R_c = 2z_R$.

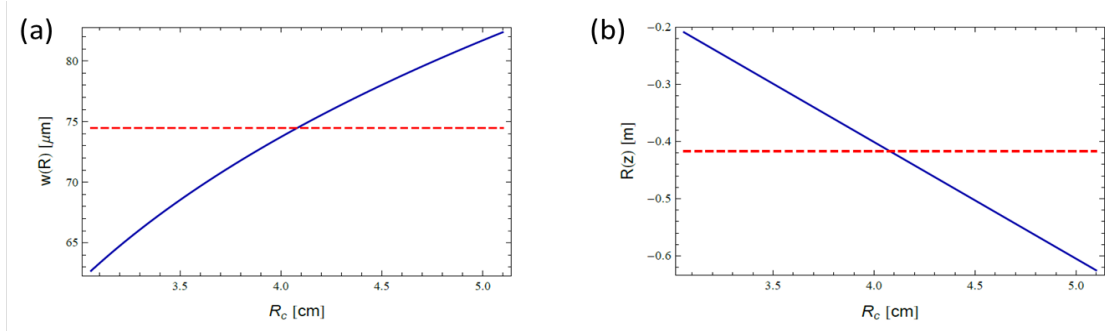


Figure 6.10: (a) Beam radius $w(d_1 + d_2)$ after a full round-trip, as a function of R_c , for $1.5z_R < R_c < 2.5z_R$. The red dashed line indicates the beam size at the input coupler, i.e. $w(0)$. (b) Radius of curvature $R(d_1 + d_2)$ after a full round-trip, as a function of R_c , for $1.5z_R < R_c < 2.5z_R$. The radius of curvature at the input coupler, $R(0)$, is marked by the dashed line.

The transfer matrix for a full round-trip is given by $M_{\text{tot}} = M_2(d_2)M_1(d_1)$,

$$M_{\text{tot}} = \begin{pmatrix} 1 & d_2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -2/R_c & 1 \end{pmatrix} \begin{pmatrix} 1 & d_1 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 - \frac{2d_2}{R_c} & d_1 - \frac{2d_1d_2}{R_c} + d_2 \\ -\frac{2}{R_c} & -\frac{2d_1}{R_c} + 1 \end{pmatrix}. \quad (6.58)$$

For this matrix, the stability criterion is

$$0 < \frac{1}{4}\text{Tr}(M_{\text{tot}}) + \frac{1}{2} < +1, \quad (6.59)$$

where

$$\begin{aligned} & \frac{1}{4}\text{Tr}(M_{\text{tot}}) \\ &= \frac{1}{4} \left(1 - \frac{2d_2}{R_c} - \frac{2d_1}{R_c} + 1 \right) \\ &= 1 - \frac{d_1 + d_2}{2R_c}. \end{aligned} \quad (6.60)$$

The stability parameters, which satisfy $0 < g_1g_2 < 1$, are therefore given by $g_1 = 1$ and $g_2 = 1 - \frac{d_1 + d_2}{2R}$. Figure 6.11 shows a plot of g_2 against g_1 , including points (1, 0), (1, 0.5) and (1, 1). The middle point (1, 0.5) corresponds to a mirror radius of $R_c = 2z_R$.




Figure 6.11: Stability diagram of the triangular ring cavity. The areas bounded by the solid line and the coordinate axes represent stable cavity configurations. The central of the three red points corresponds to a mirror curvature of $R_c = 2z_R$.

The effect of imperfect focusing can be modelled by propagating a beam with a waist of $w(x) \neq w_{\text{opt}}$ around the cavity using the ray transfer matrices M_1 and M_2 . The beam radius and radius of curvature as a function of propagation distance z are plotted in Figures 6.12 and 6.13 for different values of the waist $w(x)$.




Figure 6.12: Variation of the beam radius $w(z)$ over four passes of the total cavity length, i.e. $0 < z < 4(d_1 + d_2)$, for different values of the beam waist at the focus of the cavity. If $w_o \neq w_{\text{opt}}$, the beam requires two round-trips to be returned to its initial size.




Figure 6.13: The corresponding radius of curvature $R(z)$ for different values of w_0 .

As can be seen, the beam radius $w(z)$ of an imperfectly focused beam oscillates between two local minima after each round-trip, only returning to its initial size after traversing the cavity twice. The greater the focal mismatch between $w(x)$ and w_{opt} , the more pronounced is the difference in beam size after successive round-trips. This mismatch leads to poor overlap of the spatial modes exiting the cavity and degrades the overall signal transmission. Optimisation of the beam waist is therefore an important requirement for efficient operation of the cavity memory.

6.10 Choice of cavity parameters

Having determined a stable cavity geometry and examined the corresponding beam parameter requirements, we will now consider the choice of cavity parameters on which the attainable level of noise suppression depends. We will therefore return to our discussion of the extinction in Section 6.8 and begin by assessing the FSR and bandwidth criteria.

6.10.1 Free-spectral range and bandwidth

In order to suppress FWM noise most effectively, the cavity FSR has to be chosen such that the anti-Stokes frequency falls exactly between two transmission peaks, when the cavity is tuned into resonance with the signal frequency. Hence, the FSR $\Delta\nu$ should satisfy the following relation

$$\Delta\nu = \frac{2 \times \Delta_{\text{HFS}}}{n + 0.5} \quad (6.61)$$

where $n \geq 0$ is the number of harmonics spanning the frequency difference between the signal and the anti-Stokes component $\Delta_{\text{sig-AS}} = 2\Delta_{\text{HFS}}$. Additionally, the bandwidth of the resonance features should be larger than the bandwidth of the pulses in order to maximise the overlap between the cavity mode and the signal field. This requirement ensures that the broad acceptance bandwidth of the memory, and hence its time-bandwidth product, is maintained. We will therefore set $\delta\nu = 1.2$ GHz. As noted in Section 6.8, the high-bandwidth requirement places a strict upper bound on the finesse and consequently limits the cavity visibility and obtainable extinction. This demonstrates the sensitive trade-off between broadband memory operation and noise suppression.

For a fixed resonance bandwidth, optimal noise extinction is achieved if the FSR is twice as long as $\Delta_{\text{sig-AS}}$, i.e. for $n = 0$, $\text{FSR} = 2 \times 2\Delta_{\text{HFS}}$. However, this corresponds to a cavity length of approximately 4 mm which would pose a significant technical challenge⁵. For the purpose of a cavity memory prototype, it is sensible to sacrifice some of the extinction in favour of increasing the cavity length to a feasible size. Table 6.1 shows sets of possible cavity parameters for $n = 0, 1, 2, 3$:

⁵The optimal cavity length may be achievable with a monolithic cell-cavity system which would necessitate a miniaturised vapour cell.

n	FSR [GHz]	l [mm]	Υ'
0	36.8	4.1	0.004
1	12.27	12.2	0.03
2	7.35	20.4	0.09
3	5.26	28.5	0.17

Table 6.1: Sets of cavity parameters for different values of FSR, with corresponding cavity lengths and effective extinction factors.

Given that the wall thickness of standard vapour cells is approximately 1 mm, a cavity length on the order of 1 cm would limit the cell length to just a few millimetres. I therefore decided to select the $n = 2$ harmonic, corresponding to a round-trip cavity length of $2 \times l = 40.8$ mm and an FSR of $\Delta\nu = 7.35$ GHz, as illustrated in Figure 6.14.

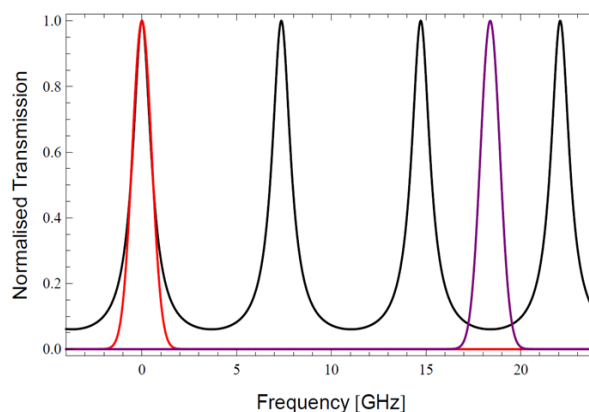


Figure 6.14: Normalised cavity spectrum for $n = 2$. While the signal field (red) is on resonance, the frequency of the anti-Stokes field (purple) falls between two transmission peaks. The frequency axis has been chosen to show the detuning relative to the signal frequency.

6.10.2 Control field coupling

Of key importance for intracavity memory operation is the ability to couple both the signal and control fields simultaneously into the cavity. This double-resonance condition can be achieved by exploiting the orthogonality of the signal and control field polarisations. A tunable uniaxial birefringent crystal placed inside the cavity can be used to bring about a

polarisation-dependent change in optical path length which shifts the resonance frequencies by an amount equal to the caesium hyperfine splitting $\Delta\nu_{\text{HFS}}$:

$$\nu_c = \nu_s + \Delta\nu_{\text{HFS}} = m \frac{c}{2nl_c} = \frac{2nl_s\nu_s}{c} \frac{c}{2n(l_s - \Delta l)} = \frac{l_s\nu_s}{l_s - \Delta l} \quad (6.62)$$

Solving for the change in optical path length gives $\Delta l = \frac{l_s \Delta\nu_{\text{HFS}}}{\nu_s + \Delta\nu_{\text{HFS}}}$.

However, due to the periodicity of the cavity resonances, path length changes corresponding to multiples of the free spectral range $\Delta\nu$ are equivalent. The set of Δl , which satisfy the above condition are therefore given by

$$\Delta l = \frac{l_s(\Delta\nu_{\text{HFS}} + m\Delta\nu)}{\nu_s + \Delta\nu_{\text{HFS}} + m\Delta\nu}, \quad (6.63)$$

where $m \in \mathbb{Z}$.




Figure 6.15: Normalised transmission spectrum of a doubly resonant cavity for orthogonal linear polarisations of the signal (solid line) and control (dashed line). The optical pumping laser (green) is partially resonant with the cavity spectrum corresponding to the control polarisation (see Section 6.11). The signal, control and anti-Stokes pulses are shown in red, blue and purple, respectively.

Figure 6.15 shows the transmission spectrum of a cavity satisfying the double-resonance condition for the signal and control field polarisations. The parameters used to plot this

spectrum are discussed in Section 6.10.3. In order to minimise spatial walk-off and absorption loss through the crystal, the smallest solution ($m = -1$) was selected, corresponding to $\Delta l \approx \frac{\lambda}{8}$. The birefringence Δn necessary to bring about the $\frac{\lambda}{8}$ -change in path length therefore guides the choice of birefringent material, given that the limited intracavity space restricts the crystal thickness to < 1 mm. The temperature-dependent change in optical path length is given by $\Delta l(T) = t_{\text{crystal}}(T)(n_o(T) - n_e(T))$ where $t_{\text{crystal}}(T)$ is the crystal thickness at temperature T . The refractive indices n_o and n_e for light polarised perpendicular and parallel to the optic axis can be found from empirical Sellmeyer equations. The temperature dependence of the material thickness $t_{\text{crystal}}(T)$ is given by the thermal expansion coefficient a . The rate of change of the birefringence with temperature can be found using the thermo-optic coefficients, dn_o/dT and dn_e/dT . The signal and control modes therefore acquire a temperature-dependent difference in path length given by

$$\Delta l(T) = t_{\text{crystal}}(T_0)e^{a(T-T_0)} \times \left[n_e(T_0) - n_o(T_0) + \left(\frac{dn_e}{dT} - \frac{dn_o}{dT} \right) \times (T - T_0) \right] \stackrel{!}{=} \lambda(n + 1/8). \quad (6.64)$$

Here we have ignored the difference in frequencies between the signal and control fields, and assumed that the signal and control fields are polarised along the principal axes of the crystal. The path length change Δl can be optimised for signal and control transmission by tuning both the angle and the temperature of the birefringent element. The temperature tunability of α -barium borate⁶, or α -BBO, is examined in Figure 6.16 which shows $t_{\text{crystal}}(T) \times \left(\frac{dn_e}{dT} - \frac{dn_o}{dT} \right) \times T$ for a crystal thickness of 0.8 mm. According to this model, a temperature change of approximately 18°C is necessary to achieve a path length difference of $\lambda/8$.

⁶The thermal expansion coefficient is given by $a = 4 \times 10^{-6}/K$ and the thermo-optic coefficients are $dn_o/dT = -9.3 \times 10^{-6}$ and $dn_e/dT = -16.6 \times 10^{-6}$ ^[305].




Figure 6.16: Change of optical path length in units of λ as a function of temperature for 0.8 mm of α -BBO.

6.10.3 Cavity mirrors

The choice of cavity bandwidth and free-spectral range constrain the value of ζ , according to equation (6.40). For a given loss parameter η , we can then determine the ideal mirror reflectivities for which the condition of critical cavity coupling is satisfied. The total intracavity round-trip transmission was measured to be $\eta^2 = 0.978$. This includes a transmissivity of 99.5% for both the vapour cell and the birefringent element and a reflectivity of 99.8% for the concave mirror. For a FWHM bandwidth of $\delta\nu = 1.2$ GHz, the optimal reflectivities of the input and output mirrors are then equal to $R_1 = |r_1|^2 = 0.605$ and $R_2 = |r_2|^2 = \frac{|r_1|^2}{\eta^2} = 0.619$, respectively. Using equation (6.40), we find that this choice of R_1 and R_2 corresponds to a cavity finesse of $F \sim 6$.

In practice, the reflectivities $R_{1/2}$ of the coupling optics are expected to differ by some amount $\epsilon_{1/2}$ from their optimal values. The sensitivity of the cavity response to changes in $\epsilon_{1/2}$ can be gauged by examining the dependence of the resonant signal transmission $I_{\text{trans,res}}$ as well as the cavity bandwidth $\delta\nu$ on the reflectivities. By way of example, Figure 6.17 shows $I_{\text{trans,res}}$ and $\delta\nu$ as a function of R_2 .

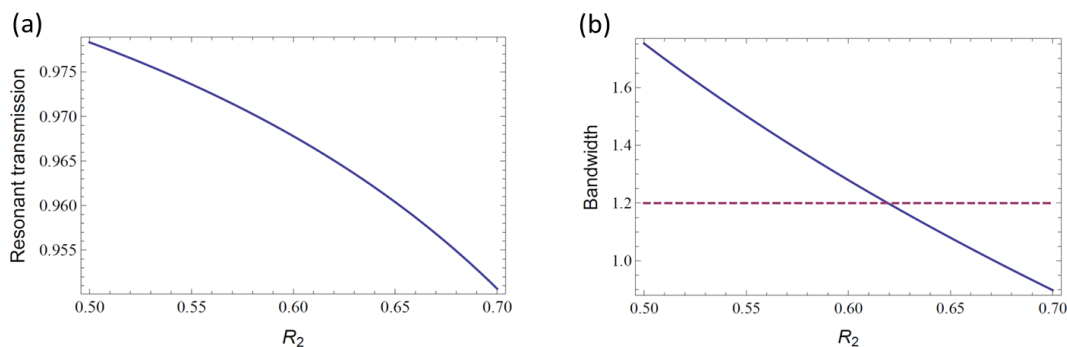


Figure 6.17: Peak transmission (a) and resonance bandwidth (b) as a function of R_2 .

While the on-resonance transmission varies only weakly with the mirror reflectivity, Figure 6.17 (b) demonstrates that the cavity bandwidth is highly sensitive to changes in the reflectivity. For instance, one finds that a 5%-change in both R_1 and R_2 results in a bandwidth variation of more than 100 MHz. The degree of precision of the mirror coatings is therefore an important design consideration. However, as we shall see in Section 7.1.2, the minimisation of $\epsilon_{1/2}$ will turn out to be a complex engineering challenge.

6.11 Optical pumping inside the cavity

In the previous section we discussed how the cavity can be made doubly resonant for both the signal and control frequencies. In order to initialise the memory medium through optical pumping, a third frequency needs to be introduced into the cavity. Fortunately, the diode laser at a frequency of $\nu_d = \nu_s - 6$ GHz is partially resonant with the cavity transmission function if prepared in the same polarisation as the control field (Figure 6.15). This is indeed a convenient choice, since polarisation filtering can be used to avoid leakage of the diode laser into the signal arm. For this purpose, the diode laser is directed into the cavity in a counter-propagating fashion via the other input port of the incoupling mirror, echoing the configuration of the free-space memory.

Achieving a sufficiently high diode laser power inside the cavity is nevertheless expected to be challenging, since resonant absorption losses are likely to preclude the benefits of strong cavity enhancement. It will therefore be necessary to ensure that the diode laser power incident on the cavity is sufficiently high to compensate for these losses.

6.12 Hänsch-Couillaud locking of the cavity

The coupling of the signal and control frequencies into the cavity is highly sensitive to perturbations of the cavity length on the sub-wavelength scale. Efficient and stable memory operation therefore necessitates active stabilisation of the cavity via a phase-sensitive locking mechanism. This requires a frequency-stable source of light whose cavity-dependent phase response can be used to derive a suitable error signal for electronic feedback on the cavity length.

In 1980 Theodor Hänsch and Bernard Couillaud proposed a locking scheme which would come to bear their name^[306]. The Hänsch-Couillaud (H-C) method provides a way of locking a cavity to the centre of a resonance feature by means of polarisation spectroscopy, without the need for any modulation techniques^[307]. It involves decomposing the incident light into two orthogonal linearly polarised components which are parallel and perpendicular to the transmission axis of an intracavity linear polariser. The reflected field of the parallel component $E_{\text{refl}}^{\parallel}$ acquires a frequency-dependent phase shift relative to the perpendicular component E_{refl}^{\perp} . The resulting elliptical polarisation of the reflected beam is analysed using a quarter-wave plate and a polarising beam splitter (PBS), as shown in Figure 6.18.




Figure 6.18: Hänsch-Couillaud locking scheme. The cavity-dependent phase shift between the reflected modes E^{\parallel} and E^{\perp} (defined with respect to the intracavity birefringence) is determined through polarisation analysis, i.e. a measurement of the third Stokes parameter S_3 .

The wave plate transforms the circularly polarised components of the elliptical beam into orthogonal linearly polarised waves which are separated into the two output ports a and b of the beam splitter. Assuming that the fast axis of the wave plate is parallel to the intracavity polariser and at 45° to the PBS, the field amplitudes measured in modes a and b are given by

$$E_{a,b} = \frac{1}{2} \begin{pmatrix} 1 & \pm 1 \\ \pm 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \begin{pmatrix} E_{\text{refl}}^{\parallel} \\ E_{\text{refl}}^{\perp} \end{pmatrix} \quad (6.65)$$

The corresponding intensities in the a and b output modes are

$$I_{a,b} \propto \left| E_{\text{refl}}^{\parallel} \pm iE_{\text{refl}}^{\perp} \right|^2 \quad (6.66)$$

The phase difference between the parallel and perpendicular components of the locking beam can be obtained by measuring the Stokes parameters $S_i = 2I_i - I_0$ where I_0 is the total intensity and $i = 1, 2, 3$ refer to the horizontal, diagonal and right-circular components, respectively^[308].

For convenience, the Stokes parameters can be represented in the (H,V) basis:

$$\begin{aligned}
 S_0 &= |E_H|^2 + |E_V|^2 \\
 S_1 &= |E_H|^2 - |E_V|^2 \\
 S_2 &= 2\Re(E_H E_V^*) \\
 S_3 &= 2\Im(E_H E_V^*)
 \end{aligned}
 \tag{6.67}$$

The locking signal used in H-C locking is given by S_3 , i.e. the intensity difference between orthogonal polarisation modes in the circular basis (\hat{L}, \hat{R}) where $\hat{L} = (\hat{H} + i\hat{V})/\sqrt{2}$. This is equivalent to measuring the phase difference between the horizontal and vertical polarisations.

The relative phase shift between the orthogonally polarised components can be expressed as a function of the change in cavity length, i.e.

$$\delta_{\parallel, \perp} = \frac{4\pi n_{e,o}(l_{\parallel, \perp} + x)\nu}{c}
 \tag{6.68}$$

where $2(l_{\parallel, \perp} + x)$ is the round-trip cavity length and x is the displacement of the concave cavity mirror. The index (\parallel, \perp) refers to the s- and p-polarised components⁷ of the locking beam with respect to the birefringent crystal. The refractive indices n_e and n_o can be determined from the Sellmeier equations for α -BBO at the wavelength of the locking laser (as we will see in Section 7.6.1, a Helium Neon laser at $\lambda = 633$ nm will be used to lock the cavity).

⁷The electric field of s-polarised light – from the German *senkrecht* – is polarised perpendicular to the plane of incidence, while p-polarised light is polarised *parallel* to the plane of incidence.




Figure 6.19: Locking signal as a function of the change in cavity length.

Figure 6.19 shows the predicted locking signal as a function of the change in cavity length, based on measurements of the reflectivity of the cavity components at 633 nm. Here we have assumed that the birefringent crystal is optimised for transmission of the signal and control frequencies, as discussed in Section 6.10.2.

As shown in a recent publication by Andrew White^[309], a more sensitive locking method for small detunings from resonance can be achieved by measuring an additional Stokes parameter. For instance, the ratio of S_3 and S_2 yields the phase difference ϕ between the H and V polarisations, while being independent of their relative intensities:

$$\tan \phi = \frac{S_3}{S_2}. \quad (6.69)$$

For the sake of experimental simplicity, the method of H-C locking was nevertheless chosen to stabilise the cavity. However, the locking set-up can be easily adapted to implement the more sophisticated and sensitive technique proposed in reference^[309].

Chapter 7

Experimental implementation of the cavity Raman memory

Based on the design considerations developed in the previous chapter, the Raman memory was realised for the first time inside a low-finesse optical cavity, enabling low-noise, broadband storage and on-demand retrieval at room temperature. In this chapter, I will discuss the experimental implementation of this novel Raman memory, including the testing of all principal components, the building of the experiment and finally, the demonstration of intracavity memory operation as well as the suppression of FWM noise.

Over the course of almost a year and a half, the cavity Raman memory has evolved from an initial idea to a mature experiment. I designed and developed this experiment in close collaboration with post-doctoral researcher Dylan Saunders. Within a year of beginning this work, Joseph Munns from the Imperial College Doctoral Training Centre joined the project. He is set to lead the next generation of cavity Raman memories during the remainder of his PhD. More recently, post-doctoral researcher Patrick Ledingham and visiting PhD student Cheng Qiu from the East China Normal University joined the ever-growing team

and provided valuable assistance with the final measurements described in the last section of this chapter. I will highlight individual contributions to the experiment directly in the following sections. Unless otherwise stated, the work presented in this chapter has been carried out by myself.

7.1 The ring cavity

7.1.1 Optical components

The component elements of the ring cavity were held in place using custom-designed mounts, allowing precise and independent control of the spatial orientation of each element. This was particularly important considering the challengingly small scale of the ring cavity geometry.

The input coupling mirror was clamped inside a cylindrical brass mount. The output mode reflected from the mirror was made optically accessible via a precisely angled tunnel drilled through the mount which was then fixed to a three-axis stage (Thorlabs NanoMax MAX313D/M).

The concave mirror was glued to a piezo-electrically-controlled ceramic ring actuator (Noliac NAC2123) and mounted to a standard LINOS mirror mount (Lees LM2-3025) which was attached to a one-dimensional translation stage. This made it possible to adjust the cavity length precisely and reversibly.

Finally, the output coupling prism was glued to the very edge of a triangular aluminium plate which itself was held by a mirror mount. Compared to the other cavity elements, this mount provided the least passive stability due to the extremely small contact area between the prism and its mount. However, this solution was necessary in order to achieve the required proximity between the input and output couplers which is on the order of a

few hundred μm . As we shall see, the method used for active locking of the cavity length is capable of compensating for any mechanical vibrations of the optical mounts. Cavity locking (see Section 7.6) is therefore indispensable for stable operation, given the current experimental realisation.

7.1.2 Coatings

One of the key requirements for realising a high-transmission cavity is the correct choice of reflectivities. As shown in Section 6.7, the condition for critical cavity coupling is satisfied if $r_1 - r_2\eta = 0$ where r_1 and r_2 are the amplitude reflection coefficients of the input and output couplers, respectively, and $1 - |\eta|^2$ is the intracavity loss for a complete round-trip. η was determined experimentally by measuring the transmission through the anti-reflection coated vapour cell¹ and the α -BBO crystal, as well as the reflectivity of the concave mirror. The combined loss factor η was then calculated from $\eta^2 = T_{\text{cell}}^2 \times T_{\text{BBO}}^2 \times R_{\text{concave}} = 0.978$.

In order to achieve these optimal reflectivities at 852nm, the input and output couplers were custom-coated by Richard Makin from the Thin Film Facility at Oxford University. One complication arose from the fact that for a given dielectric coating, the orthogonally polarised signal and control fields can experience significantly different reflectivities if the angle of incidence is not normal². This can be understood by considering the continuity condition for the tangential electric and magnetic field components of s- and p-polarised light at the interface between two materials of different refractive indices. The Fresnel equations can be used to compute the coefficients of reflection for both orthogonal polarisations. The difference in reflectivities for s- and p-polarised light were minimised by carefully choosing

¹The cell was heated using a heat gun during transmission measurements in order to prevent condensation of caesium atoms on the cell windows.

²The angle of incidence for the chosen cavity geometry is approximately 43.5° .

the number and thickness of multiple coating layers. Based on simulations performed by Richard Makin, the optical elements were successively coated with ten alternating layers of TiO_2 and SiO_2 via electron beam vapour deposition (EBVD). The thickness of these layers was measured interferometrically after each coating run. The following plot shows the reflectance at a 45° -incidence angle as a function of wavelength for s- and p-polarisations:




Figure 7.1: Reflectance of the dielectric coating at an angle of incidence of 45° for s- and p-polarised light as a function of wavelength. Experimental data is shown alongside model predictions which take into account the thickness of each coating layer. Measurements and simulations were performed by Richard Makin.

Twelve different test pieces – six each for the input and output couplers – were coated in this fashion. Each element was positioned at a slightly different distance from the central axis of the EBVD machine, resulting in a spread of reflectivities from which the most suitable could be selected. After measuring the reflectivity and transmissivity for horizontally and vertically polarised light, I chose the pair of input and output couplers whose reflectivities not only matched the condition for critical cavity coupling most closely, but also corresponded to a finesse $F \approx \pi\sqrt{R}/(1 - R)$ sufficiently low to achieve a suitably high resonance bandwidth $\delta = \Delta/F$. The reflectivities of the mirror (input coupler) and prism (output coupler) are given in Table 7.1:

	R_{in}	R_{out}
V-pol.	73.9	75.4
H-pol	56.0	55.6

Table 7.1: Measured reflectivities of input and output coupling mirrors for vertical and horizontal polarisations.

In order to achieve the best possible suppression of the anti-Stokes polarisation mode, the signal polarisation was chosen to be vertical, corresponding to the set of higher reflectivities in Table 7.1. This choice ensures that the cavity finesse – and hence the visibility of the cavity response – will be highest for the signal polarisation.

Due to the slight departure from the condition of critical cavity coupling, the signal transmission on resonance is expected to be less than predicted for optimal reflectivities, since

$$\frac{\int_{-\infty}^{\infty} d\nu I_{\text{trans}}(\nu, \zeta_{\text{actual}}) \times f(\nu - \nu_s)}{\int_{-\infty}^{\infty} d\nu f(\nu - \nu_s)} < \frac{\int_{-\infty}^{\infty} d\nu I_{\text{trans}}(\nu, \zeta_{\text{optimal}}) \times f(\nu - \nu_s)}{\int_{-\infty}^{\infty} d\nu f(\nu - \nu_s)} \quad (7.1)$$

However, given the reflectivities in Table 7.1, the extinction ratio is expected to improve from approximately 9% to 4.4%, at the expense of a reduced cavity bandwidth of ~ 0.7 GHz.

7.2 Birefringent crystal

Based on the calculations in Section 6.10.2, a $5 \times 5 \times 0.8$ mm a-cut α -BBO crystal (Newlight Photonics) was chosen to introduce the intracavity birefringence necessary to allow resonant transmission of both the signal and control polarisations.

The crystal was affixed with thermal epoxy to a thin aluminium frame (thickness ~ 1 mm) which extended into the cavity via a long arm. The temperature of the crystal was

controlled with a TEC Peltier element sandwiched between the mount and an aluminium base acting as a heat sink. The mount and base were thermally insulated to minimise heat dissipation to the environment. A Thorlabs PID temperature controller (TTC001 T-Cube) was used to actively stabilise the temperature over a range of approximately 20°C. The entire assembly was screwed to the rotating platform of a Thorlabs XYR translation stage so that the crystal axis could be angle-tuned with respect to the cavity beam path.

7.2.1 Alignment of crystal axes

The direction of the crystal's optic axis with respect to the H/V basis of a PBS was determined using the test set-up shown in Figure 7.2.




Figure 7.2: Test set-up used to determine the orientation of the crystal axes of the BBO birefringent element.

The horizontally polarised light exiting the first PBS was rotated by the first $\lambda/2$ -waveplate ($\text{HWP}_1(\theta)$) through an angle of 2θ where θ is the angle between the incident polarisation vector and the waveplate's fast axis, with refractive index $n_e < n_o$. In Jones vector notation this can be expressed as $|\Psi\rangle_1 = U_{\text{HWP},1}(\theta) |H\rangle$. If the second $\lambda/2$ -waveplate is set to the same angle θ , then $|\Psi\rangle_2 = U_{\text{HWP},2}(\theta) |\Psi\rangle_1 = U_{\text{HWP},2}(\theta) U_{\text{HWP},1}(\theta) |H\rangle = |H\rangle$ so that a projective measurement of $|\Psi\rangle_2$ on $|V\rangle$ gives zero. If we now insert the birefringent

crystal between HWP₁ and HWP₂, then the transformation becomes

$$|\Psi\rangle_2 = U_{\text{HWP},2}(\theta)U_{\text{crystal}}U_{\text{HWP},1}(\theta)|H\rangle, \quad (7.2)$$

where

$$U_{\text{crystal}} = R(\theta') \begin{pmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{pmatrix} R(-\theta') \quad (7.3)$$

and θ' is the angle between the crystal's optic axis and the $|H\rangle$ -polarisation vector. ϕ is the relative phase imparted on the polarisation components aligned with the ordinary and extraordinary axes of the crystal. If $\theta' = \theta$, $|\Psi\rangle_1$ is retarded by an amount equal to $2\pi n_e d/\lambda_0$ which can only be observed interferometrically. However, $\theta' \neq \theta$ would manifest itself as a rotation of $|\Psi\rangle_1$, resulting in a non-zero projection of $|\Psi\rangle_2$ onto $|V\rangle$.




Figure 7.3: P_V as a function of the waveplate angle θ [°]. The data was fitted with a \cos^2 -function, according to Malus's law.

Figure 7.3 shows the power P_V measured at the $|V\rangle$ -port of the second PBS as a function of the angle θ of the two waveplates. The minimum value of P_V is at $\theta \approx -1^\circ$ which suggests that the crystal axes are not perfectly aligned with respect to the frame of the beam splitters. This slight rotation will have to be compensated in order to achieve optimal suppression of the control field (see Section 7.8.1).

7.2.2 Temperature control

The same test set-up shown in Figure 7.2 was then used to measure the change in birefringence with temperature. The first waveplate was set to an angle of $\theta = 22.5^\circ$, thereby preparing the field in the diagonal polarisation, $|D\rangle = |H\rangle + |V\rangle$. The second waveplate was also set to $\theta = 22.5^\circ$ so that the degree of rotation induced by the birefringent element could be measured by monitoring the power at the output ports of the second PBS. Figure 7.4 shows the fractional power $P_H/(P_H + P_V)$ and $P_V/(P_H + P_V)$ as a function of crystal temperature.




Figure 7.4: Fractional powers $P_H/(P_H + P_V)$ (red) and $P_V/(P_H + P_V)$ (blue) as a function of crystal temperature. Error bars account for uncertainties in intensity readings. The data sets were fitted with a function of the form $a + b \cos(cT)$ with a , b and c as free parameters³.

The necessary change in path length of $\lambda/8$ for orthogonal polarisations is therefore achievable within the accessible temperature range. This agrees with the predicted dependence of the path length change on temperature, plotted in Figure 6.16.

³According to equation (6.64), the temperature-dependent phase shift imparted by the birefringent element can be approximated by the functional form $(c_1 + c_2T)$, where the constant $c_1 \gg c_2$. Using the identity $\cos(c_1 + c_2T) = \cos(c_1) \cos(c_2T) - \sin(c_1) \sin(c_2T)$, we can set $\cos(c_1) = b$ and $c_2 = c$, and absorb the second term into the constant a , since its dependence on T is weak for the range of temperatures considered.

7.3 Vapour cell

The small scale of the cavity imposes a strict upper limit on the possible length of the caesium vapour cell. This reduction in the spatial extent of the memory is partially compensated for by the mean number of passes N of the signal field through the cavity. As discussed in Section 6.3, N is given by the ratio of the cavity lifetime to half the round-trip time, i.e. $N = F/\pi$.

Given the cavity length proposed in Section 6.10.1, a vapour cell with a length of 12 mm and a diameter of 10 mm was custom-ordered from *Precision Glassblowing*. The cell was constructed from quartz with fused silica windows each of which was antireflection-coated for 852 nm on either side. The window thickness was chosen to be 1 mm so as not to compromise the structural integrity of the cell. In addition to the caesium sample, 10 Torr of neon buffer gas was added to increase the timescale of diffusion for the gaseous caesium. Compared to the previously used vapour cell, the buffer gas pressure was chosen to be lower in order to reduce the effects of pressure broadening and potentially allow memory operation at smaller detunings without increasing linear absorption. This would have the advantage of stronger Raman coupling as well as increased dispersion between the Stokes and anti-Stokes fields due to the non-collinear phase-matching condition (see Section 5.6.1).

7.3.1 Cell heating

The caesium vapour was heated to a temperature T by wrapping the cell with 50 cm of phosphor bronze Quad-TwistTM wire (AWG 36, Lakeshore cryotronics), connected to a DC voltage supply. The wire consists of two intertwined pairs of twisted leads (eight twists per inch), designed to cancel any magnetic fields induced by the ohmic heating current. In order to reduce heat dissipation to the environment, the cell was covered in high-temperature

insulation tape. Four individually calibrated ATC Semitec micro-thermistors (NTC) were used to monitor the temperature of the cell at different positions across its surface. The resistance measured by one of these thermistors was used as the input to a temperature control unit (Omega CN-243) which regulated the cell temperature via a simple relay mechanism: if the measured resistance was above a chosen setpoint value, a relay switch was closed to allow current flow through the heating wire to compensate for the drop in temperature. As soon as the setpoint resistance was reached, the switch was opened and the heating process was interrupted. Over the course of several hours, this control mechanism stabilised the temperature to within a standard deviation of 0.26 K. The resistances of the remaining three thermistors were monitored by connecting each in series with a fixed-value resistor⁴ and measuring the voltages across these potential dividers using an Arduino Mega 2560 microcontroller.

In order to prevent condensation of caesium atoms on the cell windows, a small region of a few millimetres at the top of the cell was left clear of heating wire so that a “cold spot” could be applied, similar to the previous free-space design (see Section 3.1.1). However, instead of using compressed air, which may have destabilised the cavity as a result of turbulence, a “cold finger” was constructed from a thermally insulated wire fixed to a TEC Peltier element with thermal epoxy. The temperature of the cold finger could be regulated with a DC voltage supply in order to prevent the build-up of caesium on the cell windows, while maintaining a sufficiently high vapour pressure.

The cell heating and temperature control were implemented and tested by Joseph Munns as part of his MRes degree, awarded by Imperial College London. This work also included measurements of the optical depth which I will summarise at the end of the following section.

⁴The resistance was matched to the thermistor resistance at 70°C.

7.3.2 Optical depth

For the class of ensemble-based memories, the optical depth d quantifies the strength of the light-matter interaction and is therefore a critical parameter in determining the optimal memory efficiency (see Section 2.2.2). Memory operation inside the cavity indeed hinges on achieving a sufficiently high optical depth which, enhanced by the cavity finesse, should countervail the necessary reduction in vapour cell length (by a factor of $\sim 7.5/(2 \times 1)$ compared to the free-space memory⁵).

In Section 2.2.1 we defined the optical depth in terms of the on-resonance absorption $I_{\text{out}}/I_{\text{in}} = \exp(-2d)$. Equating the exponent with the natural absorption coefficient α_N gives

$$\alpha_N = 2d \left(\frac{(\Gamma_N/2)^2}{(\nu - \nu_0)^2 + (\Gamma_N/2)^2} \right), \quad (7.4)$$

Here, ν_0 is the resonance frequency of the transition which, in the absence of additional broadening mechanisms, has a natural linewidth of Γ_N .

The on-resonance optical depth $d = -(1/2) \ln(I_{\text{out}}/I_{\text{in}})$ is not accessible experimentally, since the Lorentzian absorption profile is subject to inhomogeneous Doppler broadening. This arises from the temperature-dependent motion of the atoms relative to the beam direction which shifts the resonance frequency of each atom according to its particular velocity. Unlike homogeneous broadening, which affects all atoms equally, the Doppler-broadened absorption profile follows a Gaussian shape. As a result, the transition takes the form of a Gaussian convolved with a Lorentzian function, which is known as a Voigt profile. In addition, the bandwidth of the resonance is affected by the spectral width of the laser as well as pressure broadening due to atomic collisions. However, for simplicity we will neglect

⁵This takes into account the double-pass configuration of the ring cavity.

these contributions and only consider the dominant mechanism of Doppler broadening. The Doppler-broadened absorption profile is given by

$$\alpha_D = 2d^* \exp \left[-\frac{1}{2} \left(\frac{\nu - \nu_0}{\sigma_D} \right)^2 \right], \quad (7.5)$$

where $\sigma_D = \nu_0 \sqrt{k_B T / mc^2}$, T is the equilibrium temperature and m is the atomic mass of caesium. d^* denotes the effective optical depth which takes into account inhomogeneous broadening. Equating the integrated spectral power densities from equations (7.4) and (7.5), yields the relation

$$\frac{d^*}{d} = \sqrt{\pi \ln 2} \left(\frac{\Gamma_N}{\Gamma_D} \right), \quad (7.6)$$

where $\Gamma_D = 2\sqrt{2 \ln 2} \sigma_D$ is the FWHM of the Doppler profile.

The effective optical depth of the caesium vapour was determined by measuring the transmission of a low-power probe beam derived from an external cavity diode laser (ECDL). The ECDL frequency was scanned across the Doppler-broadened $^2S_{1/2}(F=3)^2P_{3/2}$ transition of the caesium D_2 line by steadily sweeping the angle of the piezo-electrically controlled external grating. This set-up is identical to the ECDL used for optical pumping, as described in Section 3.1.5. The frequency scan was calibrated using the hyperfine features of a Doppler-free saturated absorption spectrum which was recorded simultaneously from a caesium reference cell⁶. Figure 7.5 shows a schematic of the experimental set-up.

⁶The MatLAB code used for calibration was written by Krzysztof Kaczmarek.




Figure 7.5: Set-up for measurements of the optical depth. Doppler-free saturated absorption spectroscopy in a caesium reference cell was used to calibrate the frequency scan of the ECDL.

The absorption profile was measured with an avalanche photodiode (Thorlabs APD110A) and then fitted to the normalised transmission function for the ground-state hyperfine manifold F ,

$$T(\nu) = \exp \left\{ -2d^* \times \sum_{F'} S_{F,F'} \exp \left[-\frac{1}{2} \left(\frac{\nu - \nu_{F,F'}}{\sigma_{D,(F,F')}} \right)^2 \right] \right\}, \quad (7.7)$$

where the parameter $S_{F,F'}$ is the relative coupling strength for the transition from the ground-state manifold F to the set of excited-state manifolds F' . $\nu_{F,F'}$ is the corresponding transition frequency and $\sigma_{D,(F,F')}$ is the Doppler-broadened linewidth. A least-squares fit was performed with d^* and $\sigma_{D,(F,F')}$ as free parameters⁷. Additional fit parameters were included to take into account the frequency-dependence of the ECDL intensity as well as background noise from the detector.

Figure 7.6 shows the estimated values of d for different temperatures T , as measured with a thermistor on the outer surface of the cell. As it is non-trivial to relate the measured temperature to the actual vapour temperature inside the cell, the model fit should be treated with caution. Bearing this in mind and taking into account the cooperativity parameter of

⁷ d^* and $\sigma_{D,(F,F')}$ are both functions of temperature and therefore not independent parameters.

the cavity (see Section 6.3), the optical depth is expected to be ≥ 1000 for temperatures $T \geq 80^\circ\text{C}$. This is comparable to the optical depth measured previously in the free-space memory.




Figure 7.6: Estimated values of the optical depth d for different temperatures, obtained from the least-squares fit of the Doppler-broadened absorption profile. The red line represents the theoretically predicted temperature-dependence of the optical depth according to equation (3.1). The data was measured and analysed by Joseph Munns. Note that the discrepancy between experimental data and model prediction at high temperatures is likely to be the result of a steady-state effect: as the current through the heating wire is increased, heat is dissipated from the front and back facets of the glass cell so that the vapour temperature begins to saturate.

7.4 Experimental setup

The experimental set-up of the ring cavity is shown in Figure 7.7. The signal and control fields were prepared with the same set-up as for the free-space memory detailed in Section 3.1.4. The cavity memory simply bypassed the free-space memory, while the remainder of the original experimental architecture, including the filtering and detection stages, was maintained.

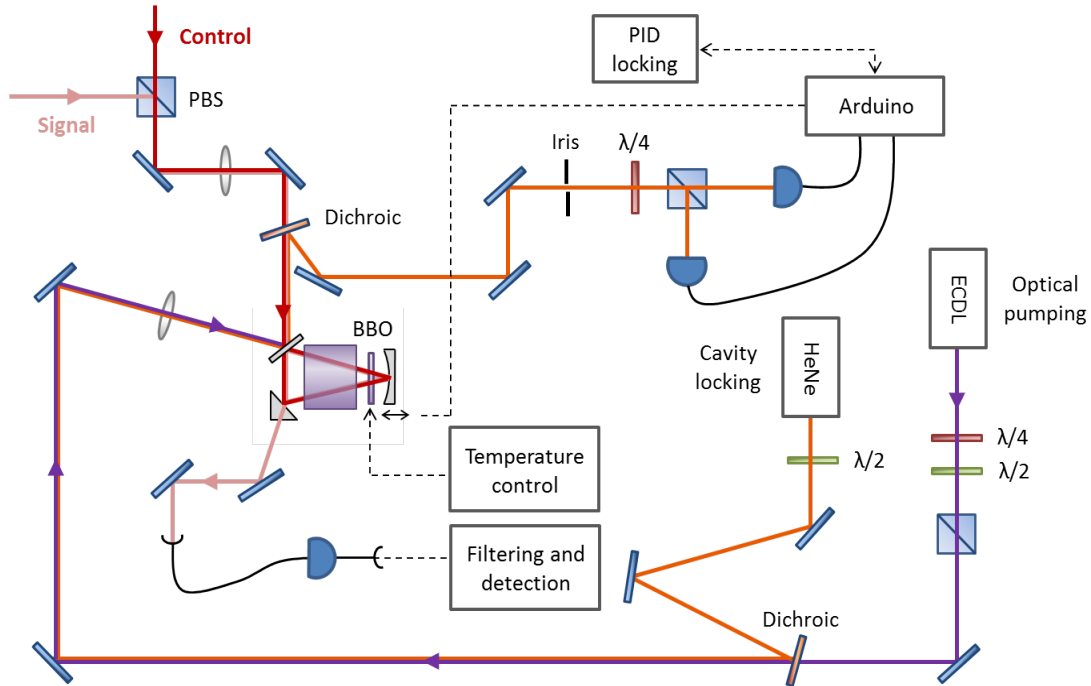


Figure 7.7: Set-up of the cavity memory. See text for details.

7.4.1 Temporal overlap of signal and control pulses

The signal and control pulses were temporally overlapped to within ± 100 ps by adjusting the relative delay of the control using a retro-reflecting stage, while monitoring the pulses on a fast photodiode with respect to the oscillator trigger. The overlap was then optimised by sending pulses of the same frequency through the signal and control arms. After recombining the orthogonally polarised pulses on a PBS, the interference visibility was then measured using a Glan-Laser polariser and a slow photodiode. In order to factor out the dependence of the visibility on spatial overlap, the pulses were coupled into a single-mode (SM) fibre after exiting the PBS. Every time the position of the retro-reflector was adjusted, the coupling efficiency of the control field into this SM-fibre was checked.

Based on these interference measurements, the temporal overlap between the signal and control pulses could be determined to within approximately 60 ps, corresponding to a sixth

of the pulse length⁸.

7.4.2 Focusing

In Section 6.9.1 it was demonstrated that the stability condition for the ring cavity requires an intracavity beam waist of $w_{\text{focus}} = 74.4 \mu\text{m}$, located half-way between the input and output couplers. This is equivalent to a FWHM of $\sqrt{2 \ln 2} w_{\text{focus}} = 87.6 \mu\text{m}$. Using a focusing lens with focal length $f = 40 \text{ cm}$, the incoming beam diameter D , defined in terms of the $1/e^2$ intensity points, can be approximated by the expression

$$D = \frac{2\lambda f}{\pi w_{\text{focus}}} \approx 3.4 \text{ mm} \quad (7.8)$$

The signal and control fields were collimated to this beam diameter using a 4:3 telescope. The beam waist w_{focus} at the focus of the 40-cm lens was then measured with a Thorlabs scanning slit beam profiler at a position corresponding to the distance between the lens and the centre of the cavity.

The focusing lens as well as one of the lenses comprising the telescope were mounted on translation stages so that the location of the beam waist could be optimised with respect to the cavity optics.

7.4.3 Piezo-electric control of the cavity length

The piezo-electric ring actuator used to control the position of the concave mirror along the cavity axis was driven with a high-power amplifier (Piezomechanik LE 150/100 EBW).

A constant DC offset voltage (0 - 150 V) could be applied to the actuator, resulting in

⁸The precision with which signal and control fields are temporally overlapped via interference measurements is limited by intensity fluctuations, the degree of spatial overlap as well as our ability to decouple the effects of spatial and temporal mode mismatch. Empirically, it has been found that the given level of precision is nonetheless sufficient to optimise the memory efficiency with respect to temporal overlap.

displacements of up to $3 \mu\text{m}$ ⁹. In addition to the manually adjustable offset voltage, the actuator could be driven with an external sinusoidal voltage signal, amplified with a variable gain of up to 30. This made it possible to sweep the cavity length continuously across multiple resonances at a frequency set by the external driving signal. In this manner the cavity spectrum could be observed during alignment (see Section 7.4.4).

The voltage response of the actuator was calibrated for light at 852 nm by setting up a standard Michelson interferometer. One arm of the interferometer was fixed with a static mirror, while the length of the other was made variable using the concave cavity mirror mounted to the actuator. The interference signal between the two interferometer arms was measured on a slow photodiode while changing the offset voltage supplied to the actuator.




Figure 7.8: Calibration of the piezo-electric ring actuator. The interference signal measured with a Michelson interferometer is plotted as a function of the DC offset voltage applied to the actuator. The error bars assume a 5% uncertainty in the interference measurements. Note that the low interference visibility – the result of imperfect alignment – was sufficient to calibrate the actuator.

The data plotted in Figure 7.8 shows that a voltage of approximately 20 V is required to displace the actuator by an amount corresponding to a complete fringe of the interference pattern. The amplitude gain setting of the piezo-electric amplifier was therefore set to provide a maximum voltage of just under 20 V. The need to limit the actuator range to a

⁹This is known as the *free stroke* of the actuator, i.e. the displacement achieved at a given input voltage in the absence of an external load.

single period follows from the fact that the laser chosen to provide the locking signal has a different wavelength than the Ti:Sa laser (see Section 7.6.2).

7.4.4 Alignment

The cavity was coarsely aligned by spatially overlapping the beams emerging from the cavity after multiple passes in order to maximise the interference fringes visible in transmission. This was done by walking the incident beam using the two mirrors immediately before the cavity, and adjusting the angle of the concave mirror. A slow photodiode was then used to further improve both the resonant transmission and the visibility of the cavity response function, whilst scanning the cavity length over several resonances. The photodiode signal was digitised via a data acquisition card (DAQ) and displayed using a LabView interface. For the signal polarisation, the optimised cavity visibility was typically found to be $\sim 85\%$. In comparison, the visibility for the orthogonal polarisation mode tended to be lower by approximately 10 – 15% as a result of the reduced cavity finesse associated with the control field polarisation (see Section 7.1.2). Note that all visibility measurements were made with the frequency of the Ti:Sa laser detuned far from both caesium resonances. This avoided the effects of the dispersive medium on the measured cavity response function. We will examine these dispersive effects more closely in Section 7.7.2.

7.5 FSR measurements

The free spectral range of the cavity was measured by recording the transmitted signal intensity, while scanning the frequency of the Ti:Sa laser over approximately 10 GHz within its mode-hop-free tuning range.




Figure 7.9: Transmission spectrum of the cavity, fitted with the transfer function I_{trans} (equation (6.39)).

Figure 7.9 shows the measured cavity response, manually fitted with the re-normalised transmission function I_{trans} which was introduced in Section 6.7. Within the experimental uncertainty, the FSR of the cavity agrees with the optimal value of 7.35 GHz. It should be noted that the data – measured directly after the cavity – displays small features between the resonance peaks which indicate that additional cavity modes were partially supported as a result of imperfect alignment. These could be suppressed by coupling the transmitted signal into a single-mode fibre.

7.6 Cavity locking

In the previous chapter we introduced the H-C method for actively stabilising the cavity length. In the following section, we will examine the experimental implementation of this locking scheme.

7.6.1 Locking signal

The locking signal was provided by a Helium Neon (HeNe) laser (Thorlabs HRS 015, 1.2 mW output power at 633 nm) whose mode of operation could be switched between intensity and

frequency stabilisation. If operated in the latter mode, the frequency is actively stabilised to within ± 2 MHz by balancing the intensity of two orthogonally polarised modes under the neon gain curve via feedback on the temperature of the laser tube^[310].

As shown in Figure 7.7, the locking beam was prepared in a superposition of orthogonal, linearly polarised components and focused into the cavity via the input mirror. The reflected output mode was separated from the incoming signal and control pulses using a dichroic mirror (reflective at 633 nm and transmissive at 852 nm). The relative phase between the polarisation components was measured by projecting the elliptically polarised output beam onto the H/V basis of a PBS. An Arduino Due microcontroller board was used to analyse the voltage signal of the two photodiodes placed at either output port of the PBS. These were connected to the analog inputs of the Arduino via a voltage divider designed to reduce the photodiode output voltage (0 - 10 V) to values within the Arduino operating range of 0 - 5 V. Determining the intensity difference between the two photodiode signals constitutes a measurement of the third Stokes parameter S_3 , as described in Section 6.12. A Python script, developed by Dylan Saunders and Joseph Munns, was used to generate the error signal relative to a given setpoint. Based on this error signal, the PID output provided kHz-feedback to the voltage applied to the piezo-electric actuator via the digitised output of the Arduino. It should be noted that the resolution of the microcontroller is limited by a 12-bit analog-to-digital converter, giving 4096 values. This discretisation ultimately limits the precision of the feedback control.

The HeNe beam incident on the cavity was aligned by maximising the visibility of the interference fringes observed in reflection, while scanning the cavity length using the piezo-electric driver. An iris positioned in front of the PBS was used as a spatial filter to increase the detected visibility. This made it possible to remove the zeroth-order reflection, i.e. the

portion of the beam which is directly reflected from the incoupling mirror without entering the cavity, from higher-order reflections which contribute to the interference visibility. In effect, post-selecting the phase-sensitive portion of the reflected light compensated for the large intracavity loss at the HeNe wavelength¹⁰ which would otherwise significantly impair the visibility.

7.6.2 Locking procedure

The locking scheme is most sensitive to perturbations if the chosen setpoint is located at the steepest part of the locking signal in Figure 6.19. In order for this setpoint to correspond to the desired cavity length, the locking signal had to be displaced laterally along the x -axis by introducing a variable phase shift between the orthogonal polarisation components of the incident HeNe beam. Initially, this was attempted by adding a variable geometric phase to the incident HeNe beam using a Berry phase set-up consisting of two quarter-waveplates on either side of a half-waveplate (see Section 4.4). However, since the intracavity loss differed for orthogonal polarisations, the change in Berry phase would cause not only a lateral shift of the locking signal, but also a displacement along the vertical axis. Instead, careful angle-tuning of the incident HeNe beam made it possible to shift the locking signal sufficiently to stabilise the cavity resonance at the signal frequency.

A more sophisticated method of locking the cavity to the correct resonance was achieved by observing the following protocol: The intensity difference at the PBS output ports, i.e. the S_3 -parameter, was first measured while scanning the cavity length in order to determine both the visibility and the mean of S_3 . The range of setpoint values, which could be manually changed by sweeping a potentiometer connected to the Arduino electronics, was

¹⁰The reflectivity of the EO3-coated concave mirror in particular dropped from $\sim 99.8\%$ at 852 nm to $\sim 20\%$ at 633 nm.

then defined to be limited by the measured visibility of S_3 . The setpoint was subsequently adjusted to the mean of the locking signal, corresponding to the steepest part of the slope.

The cavity was locked to the setpoint by replacing the sinusoidal driving signal to the piezo-electric amplifier with the PID output. Since the wavelength of the HeNe locking beam is not equal to the wavelength of the signal and control pulses, the periodicity of the cavity response differed by the ratio of the wavelengths, i.e. $633/852 \approx 3/4$. For a fixed setpoint, different locking periods therefore correspond to different levels of the signal transmission function. In order to ensure that the cavity resonance was set to the correct frequency, the signal intensity transmitted through the cavity was monitored while sweeping the DC offset voltage. If the change in offset was sufficiently large, the cavity response would be pushed into the next locking period. This process was repeated until a locking period was found for which the pre-determined setpoint corresponded to resonant signal transmission at 852 nm. The amount of signal transmission could then be optimised by slightly adjusting the setpoint via the potentiometer screw.

Large disturbances can, however, perturb the cavity length by more than one period of the locking beam, thereby stabilising the cavity towards a reduced level of signal transmission. Active stabilisation at the correct resonance is therefore only guaranteed for relatively small perturbations of the cavity. In case of larger perturbations, the offset voltage had to be adjusted manually to return the locked cavity to the signal resonance.

7.6.3 Cavity stability

The stability of the cavity lock was tested by recording the transmitted intensity of the Ti:Sa signal over the course of 45 minutes, as plotted in Figure 7.10.




Figure 7.10: Transmitted signal intensity recorded over 45 minutes with the cavity locked (dark blue) and unlocked (light blue).

In order to simulate normal operating conditions, the vapour cell was heated to 70°C and the Ti:Sa was frequency-stabilised to within a standard deviation of ~ 100 MHz. Without engaging the lock, the cavity response drifted over an entire free spectral range on timescales of tens of minutes which is indicative of temperature-induced fluctuations. In contrast, once the cavity was locked, the transmitted signal remained approximately constant, with small fluctuations about the mean corresponding to a frequency stability with a standard deviation of approximately 50 MHz. The sharp features in the trace for the locked cavity, e.g. at $t \approx 23$ min. and $t \approx 37.5$ min., were the result of knocking the optical table, demonstrating the robustness of the lock against δ -function-type disturbances.

In order to determine the typical frequencies of mechanical noise, the intensity transmitted through a mock cavity setup was measured as a function of time, following a sharp knock against the table. The data was then Fourier-transformed and compared to similar measurements taken in the absence of δ -function disturbances. These measurements demonstrated that mechanical vibrations typically occur at low frequencies of a few hundred Hz and are not present above 4 kHz¹¹. The timescale of these vibrations sets a lower limit on

¹¹The data was measured and analysed by Joseph Munns.

the speed of the feedback mechanism necessary to counteract mechanical noise.

7.7 Control field coupling

In order to render the cavity simultaneously resonant for both the signal and control fields the method of birefringence control was developed in Section 6.10.2. This method relies on a cavity response which is non-degenerate with respect to orthogonal polarisations. For the sake of completeness, I will also introduce a second method which is based on the careful control of the temperature-dependent dispersion of the caesium vapour. This second method allows for a degenerate cavity response and was developed in order to be able to increase the operating temperature of the memory.

7.7.1 Non-degenerate cavity response

As discussed in Section 6.10.2, the temperature of the intracavity birefringent crystal can be tuned such that the cavity becomes resonant for both the signal and control frequencies. A convenient method of determining the optimal temperature involves the use of a tunable CW diode laser (Toptica DL pro, < 1 MHz linewidth)¹². The frequency of this diode laser can be scanned across a mode-hop-free tuning range on the order of 20 GHz using a digital laser controller (Toptica DLC pro). Doppler-free absorption spectroscopy in a caesium reference cell allowed calibration of the frequency scan by observing the two resonances of the caesium D_2 line, separated by the hyperfine splitting ($\Delta_{\text{HFS}} = 9.2$ GHz). A $\lambda/2$ -waveplate in front of the cavity was used to rotate the polarisation of the incident field between the vertical and the horizontal, corresponding to the signal and control field polarisation, respectively. After locking the cavity to the signal frequency, the DL pro laser was directed into the cavity

¹²The DL pro laser was set up by Cheng Qiu.

while scanning its frequency across both hyperfine transitions. The polarisation was then set to diagonal so that the cavity transmission spectrum for both H - and V - polarisations could be observed simultaneously, as shown in Figure 7.11.




Figure 7.11: Cavity transmission spectra for signal and control field polarisations measured by scanning the frequency of a tunable diode laser. The temperature of the BBO crystal was set to 36°C . Note that the high-frequency intensity fluctuations on the order of several hundred MHz are due to etaloning effects.

By changing the crystal temperature, the frequency difference between resonant peaks for H and V could be optimised to give $|\nu_{\text{res},V} - \nu_{\text{res},H}| = \Delta_{\text{HFS}} - \Delta\nu \approx 1.85 \text{ GHz}$, where $\Delta\nu$ is the cavity free spectral range. This condition was met at a crystal temperature of $T = 36^\circ\text{C}$.

7.7.2 Degenerate cavity response

In our treatment of the cavity response function, we have so far not accounted for the frequency-dependent susceptibility of the caesium ensemble. However, at higher vapour temperatures ($T > 70^\circ\text{C}$), the cavity response becomes increasingly sensitive to the dispersive and absorptive properties of the Doppler-broadened ensemble. It is in fact possible to exploit this frequency dependence in order to achieve a cavity response which is resonant at the signal and control frequencies and anti-resonant at the anti-Stokes frequency, independ-

ent of polarisation. This alternative method of tuning the cavity relies on careful adjustment of the caesium temperature, rather than the intracavity birefringence.

We can include the frequency-dependent atomic response in our description of the cavity transmission function by replacing the parameter ζ in equation (6.39) with the complex term $\zeta = |\zeta(\omega)|e^{i\phi(\omega)}$. Denoting the round-trip beam path inside the caesium vapour with l_{Cs} and the remaining free-space beam path with $l_{\text{free}} = 2l - l_{\text{Cs}}$, we can write

$$\zeta(\omega) = \zeta_0 \exp \left[\frac{i\omega}{c} (n(\omega)l_{\text{Cs}} + l_{\text{free}}) \right], \quad (7.9)$$

where the complex refractive index, separated into real and imaginary parts, is given by $n(\omega) = n'(\omega) + in''(\omega)$. The absorptive and dispersive components are therefore given by

$$\begin{aligned} |\zeta(\omega)| &= \zeta_0 \exp \left[-\frac{\omega}{c} n''(\omega) l_{\text{Cs}} \right] \quad \text{and} \\ \phi(\omega) &= \frac{\omega}{c} (n'(\omega) l_{\text{Cs}} + l_{\text{free}}), \end{aligned} \quad (7.10)$$

respectively. The refractive index is related to the total linear susceptibility $\chi = \sum_{i,f} \chi_{i,f}$ according to $n = \sqrt{1 + \chi} \approx 1 + \chi/2$. The real and imaginary parts are therefore given by $n' = 1 + \chi'/2$ and $n'' = \chi''/2$. As described in reference^[311] and^[312], the contribution to the susceptibility associated with the dipole transition $|i\rangle \rightarrow |f\rangle$ is given by

$$\chi_{if}(\Delta) = f_{if}^2 \frac{d_{if}^2 N_i}{\hbar \epsilon_0} s(\Delta_{if}), \quad (7.11)$$

where $\Delta = \omega - \omega_{if}$ is the detuning from resonance, f_{if} is the transition strength factor, $d_{if} = \langle i | \mathbf{e} \mathbf{r} | f \rangle$ is the dipole matrix element and N_i is the atomic number density of the initial state $|i\rangle$. The line shape function $s(\Delta)$ can be found from the convolution of the

homogeneous atomic lineshape $f(\Delta)$ and the Gaussian velocity distribution $g(v)$ of the Doppler-broadened ensemble, i.e.

$$s(\Delta) = \int_{-\infty}^{\infty} f(\Delta - \omega v/c)g(v)dv, \quad (7.12)$$

where $f(\Delta) = i(\Gamma/2 - i\Delta)^{-1}$ and $g(v) = (u\sqrt{\pi})^{-1} \exp[-(v/u)^2]$ with $u = \sqrt{2k_B T/m}$. Here, the FWHM Γ of the Lorentzian lineshape takes into account the pressure-broadened natural linewidth^{[216] [251]}.

The lineshape function $s(\Delta)$ is related to the complex error function^[313] which is difficult to evaluate numerically^{[311] [314]}. In the limit of large detunings from resonance, i.e. $\Delta \gg \Gamma_D$ where Γ_D is the Doppler width, homogeneous broadening dominates the lineshape¹³. Applying this Lorentzian approximation, Siddons *et al.* demonstrated that $s(\Delta) = f(\Delta)$. The on-resonance optical depth $d = (\omega_0 l/c) n''$ can then be found by substituting $s(0) = 2i/\Gamma$ into the expression for $n'' = \chi''/2$,

$$d = \frac{\omega_0 l}{c} \times \frac{d_{if}^2 N_i}{\hbar \epsilon_0 \Gamma}. \quad (7.13)$$

so that the total susceptibility in terms of d is given by $\chi = dc\Gamma s(\Delta)/(\omega_0 l)$. Hence, in the far-off-resonant limit, the dispersive and absorptive components scale as

$$\begin{aligned} \text{Dispersion :} & \quad \lim_{\Delta \gg \Gamma_D, \Gamma} \left(\frac{\omega l}{c} n' \right) = \frac{\omega l}{c} - \frac{\omega}{\omega_0} \frac{d\Gamma}{2\Delta} \\ \text{Absorption :} & \quad \lim_{\Delta \gg \Gamma_D, \Gamma} \left(\frac{\omega l}{c} n'' \right) = \frac{\omega}{\omega_0} \frac{d\Gamma^2}{4\Delta^2}. \end{aligned} \quad (7.14)$$

Based on this analysis, Joseph Munns has developed a model for the degenerate cavity

¹³Siddons *et al.* interpret this result by noting that the Gaussian lineshape describing inhomogeneous Doppler broadening falls off exponentially with Δ ^[311].

response with which the optical pumping efficiency can be estimated from experimental data¹⁴.




Figure 7.12: Frequency scan of the degenerate cavity response, demonstrating resonance at the signal and control frequencies and anti-resonance at the anti-Stokes frequency (dotted lines). The two caesium hyperfine resonances at $\Delta = 0$ GHz and $\Delta = -9.2$ GHz are indicated with solid black lines. The experimentally measured cavity transmission function (blue line) was fitted with the above model (orange line) by Joseph Munns. The population inversion for this data set was estimated to be $w = p_1 - p_3 = 0.65$ at a caesium vapour temperature of $\sim 80^\circ\text{C}$. Note that the cavity was locked to the signal frequency.

Figure 7.12 shows a typical frequency scan of the cavity transmission function, demonstrating that the degenerate resonance condition can be satisfied via dispersion control. While the numerical simulation of the cavity response is a challenging problem, this work promises to provide a direct method of probing the temporal dynamics of the cavity fields.

7.8 Memory operation

Having ensured that the cavity response function fully satisfies the resonance condition, we are now ready to investigate the memory performance and – most importantly – the level of noise suppression relative to the retrieved signal. Before presenting the results of this

¹⁴The pumping efficiency can be obtained by considering the dependence of $n(\omega)$ on the initial population imbalance between the two hyperfine ground states.

investigation, I will briefly outline the methods of signal filtering and detection which are largely identical to the system used in previous memory experiments.

7.8.1 Control field suppression

In order to achieve a sufficient degree of signal filtering, it was necessary to make a slight modification to the original set-up. As discussed in Section 7.2.1, the axes of the birefringent crystal are misaligned by a small amount on the order of 1° relative to the PBS basis which defines the signal and control polarisations. With every pass of the crystal, the fields therefore acquire a small rotation which reduces the amount of control field extinction possible through polarisation filtering. Ideally, in order to achieve a complete reversal of this polarisation rotation, the signal and control fields would have to be transmitted through a second, identical cavity. However, for reasons of practicality, a $\lambda/2$ - and $\lambda/4$ -waveplate placed immediately after the cavity were instead used to partially compensate for the polarisation rotation. These waveplates were adjusted using motorised rotation mounts so that control field leakage through the filtering set-up could be minimised without inadvertently disrupting the sensitive cavity lock. Following this method, a typical polarisation extinction of approximately 4×10^3 was achievable.

The amount of control field suppression was further enhanced by adding three Fabry-Perot etalons to the frequency filtering stage. As we shall see below, minimising control field leakage is an important prerequisite for assessing the cavity performance in terms of noise suppression.

7.8.2 Signal detection and data acquisition

All measurements presented in the following sections were made with the same *Perkin Elmer* APD as used previously for the free-space memory. The arrival time of detection events relative to a trigger signal was recorded using a time-to-digital converter (*qutools* quTAU TDC) with a time bin resolution of 81 ps and a software interface developed by postdoctoral researcher Andreas Eckstein.

7.8.3 Storage of weak coherent states

Memory operation inside the cavity was demonstrated using weak coherent input states attenuated to approximately 20 photons per pulse. Figure 7.13 shows arrival-time histograms of the input signal (green), the retrieved signal after 12.5 ns (blue) and the noise present when the signal is blocked (red), as recorded with the TDC. The temporal shape of these histograms depends on the convolution of the pulse line shape and the transmission function of the cavity.




Figure 7.13: Storage and retrieval of weak coherent states of approximately 20 photons per pulse.

The memory efficiency can be evaluated from the rates of detection events $c_l^{\text{in/out}}$ in the read-in and read-out time bins with $l \in \{scd, sd, cd, d\}$ denoting the combination of signal,

control and diode (pump) fields present. In accordance with equation (5.3) in Section 5.3.2, the total memory efficiency for storage followed by retrieval is given by

$$\eta_{\text{tot}} = \frac{c_{\text{scd}}^{\text{out}} - c_{\text{cd}}^{\text{out}} - c_{\text{sd}}^{\text{out}}}{c_{\text{sd}}^{\text{in}} - c_{\text{d}}^{\text{in}}}, \quad (7.15)$$

where we have set $c_{\text{d}}^{\text{in/out}} = 0$, as the diode laser was switched off during storage events by periodically triggering the AOM-switch (see Section 3.1.5). Equation (7.15) assumes that the number of counts in the signal mode can be determined by subtracting the noise counts $c_{\text{cd}}^{\text{in/out}}$ in each time bin. It is worth clarifying this assumption: The measured detection events consist of contributions from the signal or Stokes mode, the anti-Stokes mode and the control mode. The latter contribution is the result of residual leakage of the control field through the filtering stage. We have implicitly assumed that the amount of control leakage is independent of the presence of the signal field, i.e. these leakage counts cancel in the expression $c_{\text{scd}}^{\text{out}} - c_{\text{cd}}^{\text{out}}$.

Using equation (7.15) and integrating the pulse traces in Figure 7.13 within the range indicated by the shaded areas, a total memory efficiency of $\eta_{\text{tot}} = 9.6 \pm 0.1\%$ was determined. This result represents the first experimental demonstration of Raman memory operation inside an optical cavity.

Before turning our attention to the matter of noise suppression, we will first examine the memory performance itself. Figure 7.14 (a) shows the memory lifetime measured by delaying the read-out control pulse with respect to the storage event. As expected, the memory efficiency decays exponentially with storage time, in accordance with magnetic dephasing of the spin wave coherence. However, in comparison to the free-space memory, the lifetime of the cavity memory is reduced from $1.5 \mu\text{s}$ to $94.6 \pm 7.1 \text{ ns}$ due to the absence

of magnetic shielding. Early studies of the Raman memory in free space – prior to installing μ -metal shielding – demonstrated a characteristic $1/e$ decay time of ~ 150 ns^[241]. The discrepancy between this measurement and the performance of the cavity memory can be explained by the presence of stronger magnetic fields in close proximity to the vapour cell. In order to extend the memory lifetime, it will therefore be necessary to surround the cell with appropriate magnetic shielding.




Figure 7.14: (a) Decay of memory efficiency with storage time. (b) Memory efficiency as a function of control field power.

One of the key features distinguishing the cavity implementation from free-space memories is the finesse-dependent enhancement of the coupling strength, parametrised by the cooperativity parameter D (see Section 6.3). A direct comparison of the Raman coupling constants for the free-space and cavity memories has to take into account a number of additional factors. These include the caesium number density, the ensemble length and changes to the peak Rabi frequency of the control field $\Omega_{\max} \propto \sqrt{\mathcal{E}_c/\mathcal{A}}$ as a result of intracavity transmission losses, resonant enhancement and confinement of the spatial mode volume.

Experimentally, many of these parameters are difficult to access independently without disturbing the frequency response of the cavity. In particular, it is challenging to measure

the memory efficiency as a function of caesium number density (vapour temperature) and optical pumping efficiency, since changes in these parameters can affect the cavity resonance condition. We will therefore restrict our initial investigation to a qualitative study of the dependence of memory efficiency on control field power¹⁵, as plotted in Figure 7.14 (b). At low control field powers the efficiency is expected to scale quadratically with power, as confirmed by previous measurements in free space (see Section 5.5.3). With increasing control field power, however, the A.C. Stark shift begins to perturb the Raman resonance, thus limiting the attainable memory efficiency. If the control Rabi frequency is increased further, the efficiency reaches a maximum value for a given optical depth. This effect is evident in Figure 7.14 (b) which shows the transition from quadratic scaling in the low-power regime to a saturated efficiency in the high-power regime. The optimal Rabi frequency is therefore within the range of available control field powers of $P_c \lesssim 300 \mu\text{W}$. This suggests that resonant enhancement and the tight spatial mode confinement inside the cavity compensate for coupling and transmission losses which limit the available control field power to less than half the value typical for free-space memory experiments.

For future memory realisations one could consider carving the control pulse in order to counteract the A.C. Stark shift and extend the favourable scaling of memory efficiency with control field power to higher values.

7.9 Noise suppression

The build-up of noise counts in successive time bins is a characteristic signature of four-wave mixing when a train of control pulses is incident on the memory (see Section 5.5.1).

¹⁵The power was measured before the single-mode fibre which couples the control field into the cavity. The fibre-coupling efficiency was measured to be 80%.

In the case of the free-space memory, noise counts more than quadrupled over the first four time bins, as shown in Figure 7.15 (a). This behaviour contrasts strongly with an equivalent measurement performed with the cavity Raman memory, as plotted in Figure 7.15 (b). Here, the noise increases by a factor of ~ 1.2 relative to the first read-out bin. The dramatically reduced build-up of noise is a compelling first qualitative observation of FWM noise suppression afforded by the cavity.

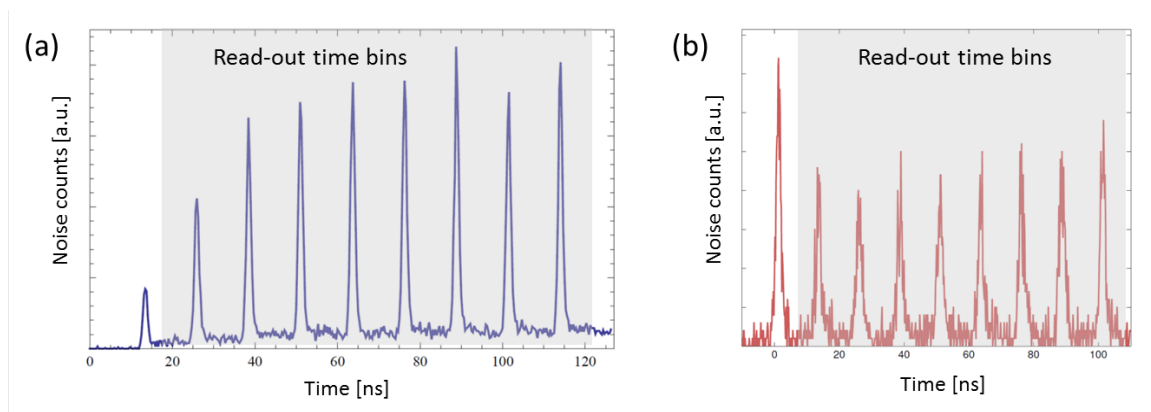


Figure 7.15: Build-up of noise counts in successive time bins for the Raman memory implemented in (a) free-space memory and (b) inside the cavity.

A noticeable feature in Figure 7.15 (b) is the large number of noise counts in the read-in time bin. This can be attributed to imperfect optical pumping, while the reduction of noise counts in the first read-out bin appears to be the result of population transfer induced by the control field. It should be noted that the control Rabi frequency is on the order of several GHz and is therefore sufficiently large to account for the nanosecond timescale of the population transfer. This explanation has been corroborated by recent results which have shown that the noise level during read-in can be reduced to the read-out level by improving the optical pumping efficiency.

A quantitative comparison between the noise level present in the cavity and free-space implementations requires an estimate of the number of noise photons per pulse, i.e. $\epsilon_{\text{in/out}}$,

retaining the notation introduced in Chapter 5. This is composed of contributions from the Stokes (S), anti-Stokes (A) and control (C) modes each of which will have a different transmission efficiency T_{filter}^m with $m \in S, A, C$ between the cavity and the detector. Hence, the rate of noise counts measured in the absence of the signal is given by

$$c_{\text{in/out}} = \eta_{\text{SPCM}} \times \sum_{m \in \{S, A, C\}} \epsilon_{\text{in/out}}^m T_{\text{filter}}^m, \quad (7.16)$$

where $\eta_{\text{SPCM}} = 0.5$ is the detector efficiency which is assumed to be the same¹⁶ for all three frequency modes m .

In our previous analysis of the free-space memory, we based our estimate of the signal-to-noise ratio on the assumption that the noise counts were dominated by contributions at the Stokes frequency. This is a valid approximation, since $T_{\text{filter}}^S \gg T_{\text{filter}}^A, T_{\text{filter}}^C$. In addition, spontaneous Stokes scattering due to imperfect optical pumping could be neglected, as this is emitted isotropically over an angle of 4π and therefore strongly suppressed through spatial filtering (see Section 5.6.1). However, in the case of the cavity memory, we should take care not to presuppose the validity of these assumptions. For instance, optical pumping of the memory medium inside the cavity is significantly more challenging than in free space, as noted in Section 6.11. The noise contribution from spontaneous Stokes scattering may therefore be higher than expected from free-space measurements as a result of the potential reduction in pumping efficiency.

Given the additional challenge of achieving optimal polarisation filtering in the case of the cavity memory (see Section 7.8.1), it is important to consider leakage of the strong control field as a potentially significant source of noise counts. We can independently measure

¹⁶This is a valid assumption, since the modes S , A and C are separated in frequency by < 20 GHz.

the amount of control leakage either by blocking the cavity or by reducing the caesium vapour temperature until the FWM contribution becomes insignificant. However, both of these methods come with a caveat; changing the vapour temperature can alter the cavity response and thus change the control field transmission through the cavity, while blocking the cavity may affect polarisation as well as frequency filtering. This latter effect is due to the fact that the total transmission efficiency of the chain of Fabry-Perot etalons depends on the pulse bandwidth which is given by the convolution of the initial pulse and the cavity transfer function. Bearing these considerations in mind, it was estimated that control field leakage contributed less than 1% to the total noise counts¹⁷.

Accurately determining the transmission efficiency of the Stokes mode between the cavity and the detector is furthermore critical to estimating the level of noise suppression. This is a non-trivial measurement, since the total transmission function T_{filter}^S is sensitive to changes in the cavity lock which may affect the bandwidth of the retrieved signal. Having determined T_{filter}^S , we can estimate the number of FWM noise photons per pulse in the Stokes mode from

$$\epsilon_{\text{in/out}}^S \approx \frac{c_{\text{in/out}}^S - c_{\text{in/out}}^C}{\eta_{\text{SPCM}} T_{\text{filter}}^S}. \quad (7.17)$$

The caveats discussed above are likely to result in an underestimate of $c_{\text{in/out}}^C$ so that equation (7.17) is expected to yield an overestimate of $\epsilon_{\text{in/out}}^S$. A lower bound on the level of FWM noise suppression in the retrieved field can then be determined by taking into account the difference in total memory efficiencies between the cavity and free-space memories. We

¹⁷This estimate was obtained by blocking the cavity. In order to compensate for the change in polarisation filtering, the transmitted power was adjusted to be approximately the same as measured when the cavity was unblocked.

will therefore define the ratio \mathcal{G}_{SNR} as

$$\mathcal{G}_{\text{SNR}} = \frac{\eta_{\text{tot,cav}}}{\eta_{\text{tot,free}}} \times \frac{\epsilon_{\text{out,free}}^S}{\epsilon_{\text{out,cav}}^S}. \quad (7.18)$$

It is important to note that \mathcal{G}_{SNR} is very sensitive to changes both in the cavity response function and in the setpoint of the cavity lock, since these have a twofold impact on \mathcal{G}_{SNR} , i.e. deviations from the optimum can degrade the memory efficiency and simultaneously increase the level of noise.

The ratio \mathcal{G}_{SNR} , as defined in equation (7.18), is equivalent to the factor by which the signal-to-noise ratio is enhanced in the cavity implementation. For a memory efficiency of $\eta_{\text{tot,cav}} = 9.6 \pm 0.1\%$, the noise was determined to be $\epsilon_{\text{out,cav}}^S = (20 \pm 2) \cdot 10^{-3}$ photons per pulse in the read-out time bin. In comparison, the equivalent noise floor in the free-space memory amounted to $\epsilon_{\text{out,free}}^S = (150 \pm 50) \cdot 10^{-3}$ photons per pulse, with a total memory efficiency of $\eta_{\text{tot,free}} = 29.0 \pm 0.9\%$ for coherent-state storage (see Sections 5.3.2 and 5.3.3). This yields a signal-to-noise enhancement factor of $\mathcal{G}_{\text{SNR}} = 2.5 \pm 0.9$ which clearly demonstrates the ability of the cavity implementation to increase the SNR well beyond previously attainable levels.

In view of future cavity memory experiments with heralded single photons, it is convenient to consider an additional figure of merit^[206], i.e. the minimum number of input photons μ needed to obtain a signal-to-noise ratio of $\text{SNR} = 1$. Based on our definition of the SNR in Section 5.3.3, this number is given by $\mu = \epsilon_{\text{out}}/\eta_{\text{mem}}$. In other words, we require a heralding efficiency of $\eta_{\text{herald}} \geq \mu$ in order to detect a retrieved signal photon with an equal or higher probability than a noise photon. Measurements of the SNR for different values of the input

photon number yield a μ -value of $\mu = 0.20 \pm 0.02$ for the cavity memory¹⁸. The condition $\eta_{\text{herald}} \geq \mu$ is therefore already met with our heralded SPDC source for which $\eta_{\text{herald}} = 0.22$, as described in Section 5.3.1.

The estimate of $\epsilon_{\text{out,cav}}^S$, which enters equation (7.18), includes contributions both from FWM noise and from spontaneous Stokes scattering. In order to compare the measured level of noise suppression with our prediction of the cavity extinction ratio, it is necessary to distinguish these two contributions. We will therefore recall our earlier analysis of the FWM interaction in Section 5.5.3, in which we discussed the dependence of FWM noise on control field energy. The quadratic scaling, which is evident in the free-space data, can be similarly examined in the cavity implementation by measuring the number of noise photons per pulse for different values of the intracavity control pulse energy¹⁹.




Figure 7.16: Number of noise photons per pulse for different control pulse energies. Horizontal error bars account for fluctuations in the measured power, while vertical error bars correspond to Poissonian errors in the detected noise counts as well as the experimental uncertainty in determining the transmission efficiency.

Figure 7.16 shows that the experimental data is well-described by a fit function of the

¹⁸In the case of the free-space memory, we find $\mu = 0.5 \pm 0.2$.

¹⁹In order to convert the measured control field power into the corresponding pulse energy, we use the method described in Section 5.5.3 and also take into account the cavity transmission efficiency.

form $ax^2 + bx + c$ with a , b and c as free parameters²⁰. Here, the quadratic term describes the FWM dependence on control pulse energy, while the linear term accounts for noise arising from imperfect optical pumping as well as any residual leakage of the control field. The constant offset represents control-independent noise, in particular collision-induced fluorescence and leakage of the diode laser. We can find an estimate of the FWM noise contribution to the measurement of $\epsilon_{\text{out,cav}}^S$ by extracting the quadratic term from the fit shown in Figure 7.16. This yields a value of $\tilde{\epsilon}_{\text{out,cav}}^S(\text{FWM}) = (6 \pm 2) \cdot 10^{-3}$, where the tilde emphasises the fact that we have corrected for all noise contributions other than FWM. Relative to free-space measurements, we arrive at a tentative estimate of the FWM noise suppression factor

$$\mathcal{G}_{\text{FWM}} = \frac{\eta_{\text{tot,cav}}}{\eta_{\text{tot,free}}} \times \frac{\epsilon_{\text{out,free}}^S}{\tilde{\epsilon}_{\text{out,cav}}^S} \sim 8, \quad (7.19)$$

where we have scaled the noise by the different memory efficiencies, as in equation (7.18). In Section 6.8 we introduced the effective cavity extinction ratio Υ' as a simple figure of merit which can be directly related to the measured cavity visibility V' . Given a typical visibility of $V' = 85\%$ for the signal/Stokes polarisation, we expect $\Upsilon' = (1 - V')/(1 + V') \approx 0.08$, the inverse of which gives a cavity suppression factor of $(\Upsilon')^{-1} \sim 12$. The difference between this factor and our estimated FWM noise suppression factor \mathcal{G}_{FWM} may in part be due to experimental imperfections when tuning the cavity response function. In fact, preliminary results based on a more complete theoretical description of the cavity memory^[315] predict a FWM noise suppression factor which appears to be consistent with our experimental estimate. This model, which is currently being developed by Joshua Nunn, also predicts that the estimated suppression factor is sufficient to result in a non-classical second-order

²⁰The values of the fit parameters are $a = (0.006 \pm 0.002)/[\text{nJ}^2]$, $b = (0.009 \pm 0.003)/[\text{nJ}]$ and $c = 0.006 \pm 0.001$.

autocorrelation of $g^{(2)} < 1$ for single-photon input signals, given the previously obtained heralding efficiency.

The above discussion suggests that the existing cavity design already holds the potential for significantly higher values of the measured SNR enhancement factor \mathcal{G}_{SNR} , provided that we can sufficiently reduce the non-FWM contributions to the total noise floor. For instance, this could be achieved by further improving the optical pumping efficiency²¹.

To date, the cavity implementation has proven to be a viable route towards achieving low-noise memory operation. First qualitative observations of noise suppression have been followed by a quantitative demonstration of significant SNR enhancement relative to the free-space memory. These encouraging results lay the groundwork for future single-photon experiments and will inform subsequent modifications to the cavity design. In the next chapter, I will briefly outline some of the future plans designed to advance the cavity Raman memory beyond this first proof-of-principle demonstration towards the goal of low-noise storage and retrieval of single photons.

²¹Higher pumping efficiencies could be achieved by increasing the diode laser power and improving the coupling efficiency. It may also be possible to reduce the detrimental effects of radiation trapping by operating the caesium vapour at lower temperatures.

Conclusion

The Raman memory experiment has over the last few years steadily evolved towards the goal of low-noise, room-temperature single-photon storage and retrieval. Beginning with its theoretical conception in 2007, the Raman memory has been the subject of extensive studies with weak coherent states at the single-photon level and with coherent polarisation states. This development culminated in the first Raman memory experiment with broadband heralded single photons derived from spontaneous parametric downconversion. The results of this investigation defined my subsequent research direction; it became apparent that a prerequisite for all future quantum applications of the Raman memory would be to find a solution to the problem of FWM noise.

Overcoming this primary obstacle is a key step in establishing the Raman memory as a practicable candidate for integration into photonic QIP technologies. Combining broadband storage, high time-bandwidth products, on-demand recall capabilities and room-temperature operation, the Raman memory is particularly well-suited for temporal multiplexing tasks. Through active synchronisation of probabilistic processes, memory-based multiplexing has the potential to transform proof-of-principle demonstrations of QIP pro-

technologies into scalable technologies.

Below I will summarise the chapters of this thesis which detail my contributions to the development of the Raman memory project.

8.1 Summary

8.1.1 Free-space Raman memory

Working alongside PhD students Duncan England and Patrick Michelberger, I investigated the phase coherence of the Raman memory interaction by storing polarisation-encoded information in a dual-rail, passively stable polarisation interferometer. Orthogonal polarisation components were mapped onto separate atomic ensembles and stored for up to $2.5 \mu\text{s}$. While the memory efficiency decayed, as expected, with increasing storage time, the results of a full process tomography showed that the process fidelity remained consistently above 84%. The contrasting behaviour suggests that dephasing of the stored spin wave does not degrade the polarisation coherence of the memory. This result is an important manifestation of the robustness of high-fidelity storage against memory loss.

Having re-built the Raman memory experiment, I worked in collaboration with Patrick Michelberger to interface the Raman memory with a type-II PPKTP waveguide source of heralded GHz-bandwidth single photons which were spectrally matched to the Raman resonance condition via frequency filtering of the herald photon. By implementing feed-forward control of the memory operation, we furthermore ensured that a fundamental requirement for on-demand synchronisation tasks had been met. The memory efficiency was determined from arrival-time histograms and bench-marked against values obtained for coherent-state signals with average photon numbers similar to the heralding efficiency. For single-

photon storage and retrieval, the total efficiency amounted to $21.1\% \pm 1.9\%$, compared to $29.0\% \pm 0.9\%$ for coherent states. The discrepancy could be attributed to a residual mismatch between the spectral modes of the single photons and the write/read pulses. Four-wave mixing (FWM) noise seeded by spontaneous anti-Stokes scattering gave rise to noise counts of $(15 \pm 5) \cdot 10^{-2}$ photons per pulse in the read-out time bin. Given the measured heralding and memory efficiencies, the signal-to-noise ratio was therefore limited to $\text{SNR} = 0.3 \pm 0.1$.

A study of the photon number statistics of the retrieved fields provided further insight into the memory and noise processes. Heralded second-order autocorrelation measurements demonstrated that single-photon input signals with $g^{(2)}(0) = 0.016 \pm 0.004$ gave rise to $g^{(2)}(0) = 1.59 \pm 0.03$ in the retrieved fields due to the presence of noise. However, equivalent measurements with coherent-state input signals attenuated to the same average photon number produced $g^{(2)}$ -values during read-out which were more than three standard deviations higher than in the case of single-photon experiments. Although non-classical photon statistics could not be preserved during retrieval, a statistically significant signature of single-photon storage was nevertheless evident in these results.

The autocorrelation measurements are well-described by a coherent interaction model which takes into account the full dynamics between the optical fields and the atomic system. The close agreement between theory and experiment not only consolidates our understanding of the FWM noise process, but also allows us to refine our description of the Raman memory interaction. The suitability of Raman memories for future quantum information applications therefore critically depends on our ability to address the problem of FWM noise.

8.1.2 Cavity Raman memory

At this important juncture, the challenge of noise suppression motivated a shift from the free-space implementation of earlier experiments towards a novel, cavity-based approach. The cavity boundaries would thereby induce a re-distribution of the density of states into which the anti-Stokes noise photon could be emitted. Besides noise suppression, a cavity implementation of the Raman memory offers the additional benefits of enhanced Raman coupling strength and high mode selectivity, independent of memory efficiency.

In collaboration with Dylan Saunders, I designed and built a low-finesse triangular ring cavity which, when tuned to the signal frequency, would be anti-resonant at the anti-Stokes frequency. I used Gaussian beam propagation to determine the parameters necessary to meet the cavity stability condition and implemented the method of birefringence control to render the cavity simultaneously resonant for both signal and control frequencies. In addition, active stabilisation of the cavity length was achieved by employing the Hänsch-Couillaud method of locking to a polarisation-dependent phase.

Using weak coherent input states, low-noise operation of the Raman memory was demonstrated for the first time inside an optical cavity, with typical memory efficiencies of $\sim 10\%$. I investigated the performance of the cavity memory in terms of the memory lifetime and the dependence of memory efficiency on control field power, showing that the memory could be operated at lower powers as a result of cavity enhancement effects. First evidence of FWM noise suppression was observed in the dramatically reduced build-up of noise counts compared to previous memory experiments in free space. Moreover, I conducted a quantitative comparison between the estimated number of noise photons in the free-space and cavity implementations and demonstrated a significant enhancement of the signal-to-noise ratio by a factor of ~ 2.5 . By identifying the FWM contribution to the measured noise

floor, I furthermore obtained an estimate of the level of FWM noise suppression possible with the current cavity design. These latest experimental results were obtained with assistance from Dylan Saunders, Joseph Munns, Patrick Ledingham and Cheng Qiu. While this is an ongoing field of work, the present investigation supports the conclusion that the cavity Raman memory offers a promising route towards future single-photon storage experiments.

8.2 Outlook

8.2.1 Single-photon storage with the cavity Raman memory

The cavity Raman memory has so far been tested and characterised with weak coherent states, demonstrating that reliable memory operation can be combined with suppression of FWM noise relative to the resonantly enhanced signal field. Having established the cavity principle as a viable option and gained insight into the experimental subtleties, we can now consider the potential ways in which the cavity design can be improved further. Increasing the current level of noise suppression will indeed be key to ensuring that non-classical photon number statistics are optimally preserved during single-photon storage. Since this is ultimately limited by the cavity visibility, the next generation of cavity memories could be designed to have a higher finesse. However, the increase in finesse and noise extinction comes at the cost of reduced memory bandwidth, as discussed in Section 6.10.1.

In addition, one could consider shortening the cavity length in order to optimise the solution to the condition imposed on the free-spectral range in equation (6.61). In the long term, a miniaturised, monolithic cell-cavity system would be ideally suited for integration into existing photonic architectures.

8.2.2 Multiphoton rate enhancement

An important example of temporal multiplexing is the synchronisation of probabilistic heralded photon sources. As discussed in Section 1.2.1, the time-bandwidth product of the Raman memory is sufficiently high to enhance multiphoton rates considerably beyond those achievable with unsynchronised sources. A first demonstration of this could in fact be performed with a single SPDC source: The memory is initially primed by storing a heralded SPDC photon. Upon detection of a second herald photon, the stored excitation is retrieved and made to coincide with the SPDC photon emitted directly from the source. The coincidence rate for these two simultaneous photons can then be compared to the probability of generating two photons from independent sources, which can be estimated from the square of the SPDC heralding efficiency. This proof-of-principle demonstration could then be extended to higher numbers of simultaneous photons by assembling additional Raman memories.

8.2.3 Light-matter Hong-Ou-Mandel interference

The Hong-Ou-Mandel (HOM) effect^[316] describes the quantum interference of two indistinguishable photons incident on either input port of a beam splitter. If the reflectivity of the beam splitter is 50%, the photons always emerge from the same output port due to destructive interference of the probability amplitudes associated with both photons being either reflected from, or transmitted through the beam splitter^[317]. The initial two-mode Fock state $|1, 1\rangle$ is thereby converted into the coherent superposition $(|2, 0\rangle - |0, 2\rangle) / \sqrt{2}$, i.e. a two-photon NOON state.

Drawing upon the unique beam splitter property of the Raman memory, an intriguingly novel variant of this phenomenon could be considered. In order to mirror the conditions

of traditional HOM interference, we will assume storage and retrieval efficiencies of 50% so that the Raman memory is equivalent to a 50:50 beam splitter for the transmitted and stored excitations. However, even if the actual efficiencies differ from these ideal values, non-classical interference should still be possible, albeit with reduced visibility.

The experimental procedure for detecting light-matter HOM interference may be as follows: Conditional on the success of a preceding storage event, a second single photon is directed into the memory, together with a control pulse. Provided that both signal photons are indistinguishable, one should be able to observe HOM-type interference between optical and material modes. Following the same reasoning as above, two possible outcomes are expected to occur with equal probability. Either the first photon is retrieved and the second photon is transmitted through the memory, or the first photon remains stored and the second photon is mapped onto another spin wave coherence. In terms of the signal amplitude S and the spin wave B , this bunching effect can be expressed as $|B_1S_2\rangle \rightarrow (|S_1S_2\rangle - |B_1B_2\rangle) / \sqrt{2}$. Here, subscripts 1 and 2 refer to the first and second optical/spin wave excitations, respectively. The two outcomes could be distinguished by initiating read-out of the memory with a further control pulse and measuring the heralded coincidence counts in the storage and retrieval time-bins.

Demonstration of light-matter HOM interference would not only represent a compelling experiment in itself, but also irrefutably assert the quantum nature of the Raman memory.

*Stands at the sea,
wonders at wondering: I
a universe of atoms
an atom in the universe.*

Richard P. Feynman

Bibliography

- [1] Bell, J. S. On the Einstein Podolsky Rosen Paradox. *Physics* **1**, 195–200 (1964).
- [2] Gisin, N., Ribordy, G., Tittel, W. & Zbinden, H. Quantum cryptography. *Reviews of Modern Physics* **74** (2002).
- [3] Nielsen, M. A. & Chuang, I. L. *Quantum computation and quantum information* (Cambridge University Press, 2000).
- [4] O’Brien, J. I., Furusawa, A. & Vučković, J. Photonic quantum technologies. *Nature Photonics* **3**, 687 (2009).
- [5] Dousse, A. *et al.* Ultrabright source of entangled photon pairs. *Nature* **466**, 217–220 (2010).
- [6] Matsuda, N. *et al.* A monolithically integrated polarization entangled photon pair source on a silicon chip. *Sci. Rep.* **2** (2012).
- [7] Bruno, N., Zambrini Cruzeiro, E., Martin, A. & Thew, R. Simple, pulsed, polarization entangled photon pair source. *Optics Communications* **327**, 3–6 (2014).
- [8] Politi, A., Matthews, J. C. F., Thompson, M. & O’Brien, J. Integrated quantum photonics. *IEEE Journal of Selected Topics in Quantum Electronics* **15**, 1673–1684 (2009).
- [9] Thompson, M., Politi, A., Matthews, J. & O’Brien, J. Integrated waveguide circuits for optical quantum computing. *IET Circuits, Devices & Systems* **5** (2011).
- [10] Metcalf, B. J. *et al.* Multiphoton quantum interference in a multiport integrated photonic device. *Nature Communications* **4** (2013).
- [11] Sprengers, J. P. *et al.* Waveguide superconducting single-photon detectors for integrated quantum photonic circuits. *Applied Physics Letters* **99**, 181110 (2011).
- [12] Pernice, W. H. P., Schuck, C., Minaeva, O., Li, M., Goltsman, G. N., Sergienko, A. V. & Tang, X., H. High-speed and high-efficiency travelling wave single-photon detectors embedded in nanophotonic circuits. *Nature Communications* **3** (2012).
- [13] Calkins, B. *et al.* High quantum-efficiency photon-number-resolving detector for photonic on-chip information processing. *Optics Express* **21**, 22657–22670 (2013).
- [14] Wootters, W. K. & Zurek, W. H. A single quantum cannot be cloned. *Nature* **299**, 802–803 (1982).

- [15] Dieks, D. Communication by EPR devices. *Physics Letters A* **92**, 271–272 (1982).
- [16] Steane, A. M. Error Correcting Codes in Quantum Theory. *Physical Review Letters* **77** (1996).
- [17] Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Physical Review A* **52** (1995).
- [18] Varnava, M., Browne, D. E. & Rudolph, T. Loss tolerant linear optical quantum memory by measurement-based quantum computing. *New Journal of Physics* **9** (2007).
- [19] Grosshans, F. & Grangier, P. Quantum cloning and teleportation criteria for continuous quantum variables. *Physical Review A* **64** (2001).
- [20] Collins, D., Gisin, N. & de Riedmatten, H. Quantum relays for long distance quantum cryptography. *Journal of Modern Optics* **52** (2005).
- [21] Briegel, H.-J., Dur, W., Cirac, J. & Zoller, P. Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication. *Physical Review Letters* **81**, 5932–5935 (1998).
- [22] Blume-Kohout, R., Caves, C. M. & Deutsch, I. H. Climbing mount scalable: Physical resource requirements for a scalable quantum computer. *Foundations of Physics* **32**, 1641–1670 (2002).
- [23] Migdall, A. L., Branning, D. & Castelletto, S. Tailoring single-photon and multiphoton probabilities of a single-photon on-demand source. *Physical Review A* **66** (2002).
- [24] Jennewein, T., Barbieri, M. & White, A. G. Single-photon device requirements for operating linear optics quantum computing outside the post-selection basis. *Journal of Modern Optics* **58**, 276–287 (2011).
- [25] Ma, X.-S., Zotter, S., Kofler, J., Jennewein, T. & Zeilinger, A. Experimental generation of single photons via active multiplexing. *Physical Review A* **83** (2011).
- [26] Christ, A. & Silberhorn, C. Limits on the deterministic creation of pure single-photon states using parametric down-conversion. *Physical Review A* **85**, 023829 (2012).
- [27] Akopian, N. *et al.* Entangled Photon Pairs from Semiconductor Quantum Dots. *Physical Review Letters* **96**, 130501 (2006).
- [28] Young, R. J., Stevenson, R. M., Atkinson, P., Cooper, K., Ritchie, D. A. & Shields, A. J. Improved fidelity of triggered entangled photons from single quantum dots. *New Journal of Physics* **8** (2006).
- [29] Hafenbrak, R., Ulrich, S. M., Michler, P., Wang, L., Rastelli, A. & Schmidt, O. G. Triggered polarization-entangled photon pairs from a single quantum dot up to 30 K. *New Journal of Physics* **9** (2007).
- [30] Förtsch, M. *et al.* A versatile source of single photons for quantum information processing. *Nature Communications* **4** (2013).
- [31] Yuan, Z.-S. *et al.* Synchronized Independent Narrow-Band Single Photons and Efficient Generation of Photonic Entanglement. *Physical Review Letters* **98** (2007).
- [32] Pittman, T. B., Jacobs, B. C. & Franson, J. D. Single photons on pseudodemand from stored parametric down-conversion. *Physical Review A* **66**, 042303 (2002).
- [33] Mower, J. & Englund, D. Efficient generation of single and entangled photons on a silicon photonic integrated chip. *Physical Review A* **84** (2011).

- [34] Sinclair, N. *et al.* Spectral Multiplexing for Scalable Quantum Photonics using an Atomic Frequency Comb Quantum Memory and Feed-Forward Control. *Physical Review Letters* **113** (2014).
- [35] Lanyon, B. P. & Langford, N. K. Experimentally generating and tuning robust entanglement between photonic qubits. *New Journal of Physics* **11** (2009).
- [36] Wieczorek, W., Krischek, R., Kiesel, N., Michelberger, P., Tóth, G. & Weinfurter, H. Experimental Entanglement of a Six-Photon Symmetric Dicke State. *Physical Review Letters* **103** (2009).
- [37] Thomas-Peter, N. *et al.* Integrated photonic sensing. *New Journal of Physics* **13** (2011).
- [38] Kwiat, P. G., Mattle, K., Weinfurter, H., Zeilinger, A., Sergienko, A. V. & Shih, Y. New high-intensity source of polarization-entangled photon pairs. *Physical Review Letters* **75**, 4337–4341 (1995).
- [39] Cohen, O., Lundeen, J. S., Smith, B. J., Puentes, G., Mosley, P. J. & Walmsley, I. A. Tailored photon-pair generation in optical fibers. *Physical Review Letters* **102**, 123603 (2009).
- [40] Mosley, P. J., Lundeen, J. S., Smith, B. J., Wasylczyk, P., URen, A. B., Silberhorn, C. & Walmsley, I. A. Heralded generation of ultrafast single photons in pure quantum states. *Physical Review Letters* **100**, 133601 (2008).
- [41] Yao, X.-C. *et al.* Observation of eight-photon entanglement. *Nature Photonics* **6**, 225–228 (2012).
- [42] Nunn, J. *et al.* Enhancing multiphoton rates with quantum memories. *Physical Review Letters* **110**, 1–5 (2013).
- [43] Gisin, N. & Thew, R. Quantum communication. *Nature Photonics* **1**, 165–171 (2007).
- [44] Bennett, C. H. & Brassard, G. Quantum cryptography: Public key distribution and coin tossing. *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing* 175–179 (1984).
- [45] Ekert, A. Quantum cryptography based on Bells theorem. *Physical Review Letters* **67** (1991).
- [46] Shor, P. W. & Preskill, J. Simple proof of security of the BB84 quantum key distribution protocol. *Physical Review Letters* **85**, 441–444 (2000).
- [47] Gorshkov, A. V. *Novel Systems and Methods for Quantum Communication, Quantum Computation, and Quantum Simulation*. Ph.D. thesis (2009).
- [48] Zoller, P. *et al.* Quantum information processing and communication - strategic report on current status, visions and goals for research in Europe. *Eur. Phys. J. D* **36** (2005).
- [49] Bouwmeester, D., Ekert, A. & Zeilinger, A. *The Physics of Quantum Information: Quantum Cryptography, Quantum Teleportation, Quantum Computation* (Springer, New York, 2000).
- [50] Nagayama, K., Kakui, M., Matsui, M., Saitoh, I. & Chigusa, Y. Ultra-low-loss (0.1484 dB/km) pure silica core fibre and extension of transmission distance. *Electronics Letters* **38**, 1168 – 1169 (2002).
- [51] Zukowski, M., Zeilinger, A., Horne, M. A. & Ekert, A. K. Event-ready-detectors Bell experiment via entanglement swapping. *Physical Review Letters* **71** (1993).

- [52] Sangouard, N., Simon, C., de Riedmatten, H. & Gisin, N. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics* **83** (2011).
- [53] Lukin, M. D. Colloquium: Trapping and manipulating photon states in atomic ensembles. *Reviews of Modern Physics* **75** (2003).
- [54] Tittel, W., Afzelius, M., Chanelie're, T., Cone, R., Kroll, S., Moiseev, S. & Sellars, M. Photon-echo quantum memory in solid state systems. *Laser Photon. Rev.* **4** (2009).
- [55] Hammerer, K., Sørensen, A. S. & Polzik, E. S. Quantum interface between light and atomic ensembles. *Reviews of Modern Physics* **82**, 1041–1093 (2010).
- [56] Bose, S., Knight, P. L., Plenio, M. B. & Vedral, V. Proposal for Teleportation of an Atomic State via Cavity Decay. *Physical Review Letters* **83** (1999).
- [57] Cabrillo, C., Cirac, J. I., García-Fernández, P. & Zoller, P. Creation of entangled states of distant atoms by interference. *Physical Review A* **59** (1999).
- [58] Duan, L. M., Lukin, M. D., Cirac, J. I. & Zoller, P. Long-distance quantum communication with atomic ensembles and linear optics. *Nature* **414**, 413–418 (2001).
- [59] Kuzmich, A., Bowen, W. P., Boozer, A. D., Boca, A., Chou, C. W., Duan, L.-M. & Kimble, H. J. Generation of Nonclassical Photon Pairs for Scalable Quantum Communication with Atomic Ensembles. *Nature* **423** (2003).
- [60] van der Wal, C. H., Eisaman, M. D., Andre, A., Walsworth, R. L., Phillips, D. F., Zibrov, A. S. & Lukin, M. D. Atomic Memory for Correlated Photon States. *Science* **301** (2003).
- [61] Matsukevich, D. N. & Kuzmich, A. Quantum State Transfer Between Matter and Light. *Science* **306**, 663–666 (2004).
- [62] Chou, C. W., de Riedmatten, H., Felinto, D., Polaykov, S. V., van Enk, S. J. & Kimble, H. J. Measurement-induced entanglement for excitation stored in remote atomic ensembles. *Nature* **438**, 828 (2005).
- [63] Chou, C., Laurat, J., Deng, H., Choi, K. S., de Riedmatten, H., Felinto, D. & Kimble, H. J. Functional Quantum Nodes for Entanglement Distribution over Scalable Quantum Networks. *Science* **316** (2007).
- [64] Yuan, Z.-S., Chen, Y.-A., Zhao, B., Chen, S., Schmiedmayer, J. & Pan, J.-W. Experimental demonstration of a BDCZ quantum repeater node. *Nature* **454** (2008).
- [65] Kimble, H. J. The quantum internet. *Nature* **453**, 1023–1030 (2008).
- [66] Cirac, J. I., Zoller, P., Kimble, H. J. & Mabuchi, H. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Physical Review Letters* **76**, 3221–3224 (1997).
- [67] Lance, A. M., Symul, T., Bowen, W. P., Sanders, B. C. & Lam, P. K. Tripartite Quantum State Sharing. *Physical Review Letters* **92**, 177903 (2004).
- [68] Feynman, R. P. Simulating physics with computers. *International Journal of Theoretical Physics* **21**, 467–488 (1982).
- [69] Deutsch, D. Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **400**, 97–117 (1985).

- [70] Shor, P. W. Algorithms for Quantum Computation : Discrete Logarithms and Factoring. *35th Annual Symposium on Foundations of Computer Science, 1994 Proceedings* 124–134 (1994).
- [71] Shor, P. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM journal on computing* **26**, 126 (1997).
- [72] Grover, L. K. A fast quantum mechanical algorithm for database search. *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC 96* **28**, 212–219 (1996).
- [73] Grover, L. K. Quantum Mechanics Helps in Searching for a Needle in a Haystack. *Physical Review Letters* **79**, 325–328 (1997).
- [74] Lloyd, S. Universal Quantum Simulators. *Science* **273**, 1073–1078 (1996).
- [75] Cirac, J. & Zoller, P. Quantum computations with cold trapped ions. *Physical Review Letters* **74**, 4091–4094 (1995).
- [76] Monroe, C., Meekhof, D. M., King, B. E., Itano, W. M. & Wineland, D. J. Demonstration of a Fundamental Quantum Logic Gate. *Physical Review Letters* **75**, 4714–4717 (1995).
- [77] Barrett, M. D. *et al.* Deterministic quantum teleportation of atomic qubits. *Nature* **429**, 737–739 (2004).
- [78] Riebe, M. *et al.* Deterministic quantum teleportation with atoms. *Nature* **429**, 734–737 (2004).
- [79] Clarke, J. & Wilhelm, F. K. Superconducting quantum bits. *Nature* **453**, 1031–1042 (2008).
- [80] Devoret, M. H. & Schoelkopf, R. J. Superconducting circuits for quantum information: an outlook. *Science* **339**, 1169–1174 (2013).
- [81] Fraval, E., Sellars, M. J. & Longdell, J. J. Dynamic Decoherence Control of a Solid-State Nuclear-Quadrupole Qubit. *Physical Review Letters* **95**, 030506 (2005).
- [82] Freedman, S. J. & Clauser, J. F. Experimental test of local hidden-variable theories. *Physical Review Letters* **28**, 938–941 (1972).
- [83] Aspect, A., Grangier, P. & Roger, G. Experimental tests of realistic local theories via Bells theorem. *Physical Review Letters* **47**, 460–463 (1981).
- [84] O’Brien, J. L. Optical quantum computing. *Science* **318**, 1567–1570 (2007).
- [85] Ladd, T., Jelezko, F., Laflamme, R., Nakamura, Y., Monroe, C. & O’Brien, J. Quantum computers. *Nature* **464**, 45–53 (2010).
- [86] DiVincenzo, D. P. Two-bit gates are universal for quantum computation. *Physical Review A* **51**, 1015–1022 (1995).
- [87] Knill, E., Laflamme, R. & Milburn, G. J. A scheme for efficient quantum computation with linear optics. *Nature* **409**, 46–52 (2001).
- [88] Bouwmeester, D. Experimental quantum teleportation. *Nature* **390**, 575–579 (1997).
- [89] Bennett, C. H. Teleporting an unknown quantum state via dual classical and EinsteinPodolskyRosen channels. *Physical Review Letters* **70**, 1895–1899 (1993).
- [90] Gottesman, D. & Chuang, I. L. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature* **402**, 390–393 (1999).

- [91] Prevedel, R., P. Walther, Tiefenbacher, F., Bohi, P., Kaltenbaek, R., T. Jennewein & Zeiling, A. High-speed linear optics quantum computing using active feed-forward. *Nature* **445**, 65–69 (2007).
- [92] Xu, X., Bao, X. & Pan, J. Demonstration of active feedforward one-way quantum computing with photon-matter hyperentanglement. *Physical Review A* **86** (2012).
- [93] Barrett, S., Rohde, P. & Stace, T. Scalable quantum computing with atomic ensembles. *New Journal of Physics* **12**, 093032 (2010).
- [94] Jozsa, R. Fidelity for mixed quantum states. *Journal of Modern Optics* **41**, 2315–2323 (1994).
- [95] Curty, M. & Lütkenhaus, N. Intercept-resend attacks in the Bennett-Brassard 1984 quantum-key-distribution protocol with weak coherent pulses. *Physical Review A* **71**, 062301 (2005).
- [96] Felix, G. N. S. A., S. & Zbinden, H. Faint laser quantum key distribution: Eavesdropping exploiting multiphoton pulses. *Journal of Modern Optics* **48**, 2009–2021 (2001).
- [97] Massar, S. & Popescu, S. Optimal Extraction of Information from Finite Quantum Ensembles. *Physical Review Letters* **74**, 1259–1263 (1995).
- [98] Hosseini, M., Campbell, G., Sparkes, B. M., Lam, P. K. & Buchler, B. C. Unconditional room-temperature quantum memory. *Nature Physics* **7**, 794–798 (2011).
- [99] Cerf, N. J., Kruger, O., Navez, P., Werner, R. F. & Wolf, M. M. Non-Gaussian Cloning of Quantum Coherent States is Optimal. *Physical Review Letters* **95**, 070501 (2005).
- [100] Hammerer, K., Wolf, M. M., Polzik, E. S. & Cirac, J. I. Quantum Benchmark for Storage and Transmission of Coherent States. *Physical Review Letters* **94**, 150503 (2005).
- [101] Specht, H. P., Christian, N., Reiserer, A., Figueroa, E., Ritter, S. & Rempe, G. A Single-Atom Quantum Memory. *Nature* **473**, 190–193 (2011).
- [102] Gilchrist, A., Langford, N. K. & Nielsen, M. A. Distance measures to compare real and ideal quantum processes. *Physical Review A* **71**, 062310 (2005).
- [103] He, Q. Y., Reid, M. D., Giacobino, E., Cviklinski, J. & Drummond, P. D. Dynamical oscillator-cavity model for quantum memories. *Physical Review A* **79**, 022310 (2009).
- [104] Sparkes, B. M. *Development of a Three-Level Gradient Echo Memory*. Honours thesis, The Australian National University (2008).
- [105] Yoran, N. & Reznik, B. Deterministic Linear Optics Quantum Computation with Single Photon Qubits. *Physical Review Letters* **91** (2003).
- [106] Nielsen, M. A. Optical quantum computation using cluster States. *Physical Review Letters* **93**, 040503 (2004).
- [107] Varnava, M., Browne, D. E. & Rudolph, T. Loss Tolerance in One-Way Quantum Computation via Counterfactual Error Correction. *Physical Review Letters* **97** (2006).
- [108] Sparkes, B. M. *Storage and Manipulation of Optical Information Using Gradient Echo Memory in Warm Vapours and Cold Ensembles*. Ph.D. thesis (2013).
- [109] Razavi, M. P. & Lütkenhaus, N. Quantum repeaters with imperfect memories: Cost and scalability. *Physical Review A* **80**, 032301 (2009).

- [110] Collins, O. A., Jenkins, S. D., Kuzmich, A. & Kennedy, T. A. B. Multiplexed Memory-Insensitive Quantum Repeaters. *Physical Review Letters* **98**, 060502 (2007).
- [111] Fiorentino, M., Spillane, S. M., Beausoleil, R. G., Roberts, T. D., Battle, P. & Munro, W. M. Spontaneous parametric down-conversion in periodically poled KTP waveguides and bulk crystals. *Optics Express* **15**, 7479–7488 (2007).
- [112] Wolfgramm, F., Xing, X., Cer e, A., Predojevic, A., Steinberg, A. M. & Mitchel, M. W. Bright filter-free source of indistinguishable photon pairs. *Optics Express* **16**, 18145–18151 (2008).
- [113] Zhang, H. *et al.* Preparation and storage of frequencyuncorrelated entangled photons from cavity-enhanced spontaneous parametric downconversion. *Nature* **5**, 628–632 (2011).
- [114] Mosley, P. J., Christ, A., Eckstein, A. & Silberhorn, C. Direct measurement of the spatial-spectral structure of waveguided parametric down-conversion. *Physical Review Letters* **103**, 1–4 (2009).
- [115] Vasilyev, D. V., Sokolov, I. V. & Polzik, E. S. Quantum memory for images: A quantum hologram. *Physical Review A* **77**, 020302 (2008).
- [116] Grodecka-Grad, A., Zeuthen, E. & Sørensen, A. S. High-capacity spatial multimode quantum memories based on atomic ensembles. *Physical Review Letters* **109**, 1–5 (2012).
- [117] Kraus, B., Tittel, W., Gisin, N., Nilsson, M., Kröll, S. & Cirac, J. I. Quantum memory for nonstationary light fields based on controlled reversible inhomogeneous broadening. *Physical Review A* **73**, 020302 (2006).
- [118] Simon, C., de Riedmatten, H., Afzelius, M., Sangouard, N., Zbinden, H. & Gisin, N. Quantum Repeaters with Photon Pair Sources and Multimode Memories. *Physical Review Letters* **98**, 190503 (2007).
- [119] Gündogan, M., Patrick M. Ledingham, A. A., Cristiani, M. & de Riedmatten, H. Quantum Storage of a Photonic Polarization Qubit in a Solid. *Physical Review Letters* **108**, 190504 (2012).
- [120] Afzelius, M., Simon, C. & Hugues de Riedmatten, N. G. Multimode quantum memory based on atomic frequency combs. *Physical Review A* **79**, 052329 (2009).
- [121] Nunn, J., Reim, K., Lee, K. C., Lorenz, V. O., Sussman, B. J., Walmsley, I. a. & Jaksch, D. Multimode memories in atomic ensembles. *Physical Review Letters* **101**, 1–4 (2008).
- [122] Parkins, A. S., P. Marte, P. Z. & Kimble, H. J. Synthesis of arbitrary quantum states via adiabatic transfer of Zeeman coherence. *Physical Review Letters* **71**, 3095–3098 (1993).
- [123] Ritter, S. *et al.* An elementary quantum network of single atoms in optical cavities. *Nature* **484**, 195–200 (2012).
- [124] Kielpinski, D., Meyer, V., Rowe, M. A., Sackett, C. A., Itano, W. M., Monroe, C. & Wineland, D. J. A Decoherence-Free Quantum Memory Using Trapped Ions. *Science* **291**, 1013–1015 (2001).
- [125] Blinov, B. B., Moehring, D. L., Duan, L.-M. & Monroe, C. Observation of entanglement between a single trapped atom and a single photon. *Nature* **428**, 153–157 (2004).
- [126] Langer, C. *et al.* Long-Lived Qubit Memory Using Atomic Ions. *Physical Review Letters* **95**, 060502 (2005).

- [127] Togan, E. *et al.* Quantum entanglement between an optical photon and a solid-state spin qubit. *Nature* **466**, 730–734 (2010).
- [128] Shim, J. H., Niemeyer, I., Zhang, J. & Suter, D. Room-temperature high-speed nuclear-spin quantum memory in diamond. *Physical Review A - Atomic, Molecular, and Optical Physics* **87**, 1–7 (2013).
- [129] Fuchs, G. D., Burkard, G., Klimov, P. V. & Awschalom, D. D. A quantum memory intrinsic to single nitrogenvacancy centres in diamond. *Nature Physics* **7**, 789–793 (2011).
- [130] Maurer, P. C. *et al.* Room-Temperature Quantum Bit Memory Exceeding One Second. *Science* **336**, 1283–1286 (2012).
- [131] Morton, J. J. L. *et al.* Solid-state quantum memory using the ^{31}P nuclear spin. *Nature* **455**, 1085–1088 (2008).
- [132] Saeedi, K. *et al.* Room-Temperature Quantum Bit Storage Exceeding 39 Minutes Using Ionized Donors in Silicon-28. *Science* **342**, 830–833 (2013).
- [133] Li, X., Voss, P. L., Chen, J., Sharping, J. E. & Kumar., P. Storage and long-distance distribution of telecommunications-band polarization entanglement generated in an optical fiber. *Optics Letters* **30**, 1201–1203 (2005).
- [134] Landry, O., van Houwelingen, J. A. W., Beveratos, A., Zbinden, H. & Gisin, N. Quantum teleportation over the swisscom telecommunication network. *Journal of the Optical Society of America B* **24** (2007).
- [135] Pittman, T. B. & Franson, J. D. Cyclical quantum memory for photonic qubits. *Physical Review A* **66**, 062302 (2002).
- [136] Leung, P. M. & Ralph, T. C. Quantum memory scheme based on optical fibers and cavities. *Physical Review A* **74**, 022311 (2006).
- [137] Song, K. Y., Herraiez, G. & Thevenaz, M. . Observation of pulse delaying and advancement in optical fibers using stimulated brillouin scattering. *Optics Express* **13** (2005).
- [138] Okawachi, Y., Bigelow, M. S., Jay E. Sharping, Z. Z., Aaron Schweinsberg, D. J. G., Boyd, R. W. & Gaeta, A. L. Tunable all-optical delays via brillouin slow light in an optical fiber. *Physics Letters A* **94**, 153902 (2005).
- [139] Stucki, D., Gisin, N., Guinnard, O., Ribordy, G. & Zbinden, H. Quantum key distribution over 67 km with a plug&play system. *New Journal of Physics* **4**, 41 (2002).
- [140] Franson, J. & Pittman, T. No Title. *Fortschr. Phys.* **46** (1998).
- [141] Tanabe, T., Notomi, M., Kuramochi, E., Shinya, A. & Taniyama, H. Trapping and delaying photons for one nanosecond in an ultrasmall high-Q photonic-crystal nanocavity. *Nature Photonics* **1**, 49–52 (2007).
- [142] Lvovsky, A. I., Sanders, B. C. & Tittel, W. Optical quantum memory. *Nature Photonics* **3**, 706–714 (2009).
- [143] Sprague, M. R. *Raman Memory for Entanglement in Diamonds and Light Storage in Optical Fibres*. Ph.D. thesis.
- [144] Boller, K.-J., Imamoglu, A. & Harris, S. Observation of Electromagnetically Induced Transparency. *Physical Review Letters* **66**, 2593–2596 (1991).

- [145] Fleischhauer, M., Imamoglu, A. & Marangos, J. Electromagnetically induced transparency: Optics in coherent media. *Reviews of Modern Physics* **77**, 633–673 (2005).
- [146] Toll, J. S. Causality and the Dispersion Relation: Logical Foundations. *Phys. Rev.* **104**, 1760 (1956).
- [147] Hau, L. V., Harris, S. E., Dutton, Z. & Behroozi, C. H. Light speed reduction to 17 metres per second in an ultracold atomic gas. *Nature* **397**, 594–598 (1998).
- [148] Fleischhauer, M. & Lukin, M. D. Dark-State Polaritons in Electromagnetically Induced Transparency. *Physical Review Letters* **84**, 5094–5097 (2000).
- [149] Liu, C., Dutton, Z., Behroozi, C. H. & Hau, L. V. Observation of coherent optical information storage in an atomic medium using halted light pulses. *Nature* **409**, 490–493 (2001).
- [150] Phillips, D. F., Fleischhauer, A., Mair, R. L. W. & Lukin, M. Storage of Light in Atomic Vapor. *Physical Review Letters* **86**, 783–786 (2001).
- [151] Cviklinski, J., Ortalo, J., Laurat, J., Bramati, A., Pinard, M. & Giacobino, E. Reversible Quantum Interface for Tunable Single-Sideband Modulation. *Physical Review Letters* **101**, 133601 (2008).
- [152] Turukhin, A. V., Sudarshanam, V. S., Shahriar, M. S., Musser, J. A., Ham, B. S. & Hemmer, P. R. Observation of Ultraslow and Stored Light Pulses in a Solid. *Physical Review Letters* **88**, 023602 (2001).
- [153] Baldit, E., Bencheikh, K., Monnier, P., Levenson, J. A. & Rouget, V. Ultraslow Light Propagation in an Inhomogeneously Broadened Rare-Earth Ion-Doped Crystal. *Physical Review Letters* **95**, 143601 (2005).
- [154] Longdell, J. J., Fraval, E., Sellars, M. J. & Manson, N. B. Stopped light with storage times greater than one second using electromagnetically induced transparency in a solid. *Physical Review Letters* **95**, 2–5 (2005).
- [155] Heinze, G., Hubrich, C. & Halfmann, T. Stopped light and image storage by electromagnetically induced transparency up to the regime of one minute. *Physical Review Letters* **111**, 1–5 (2013).
- [156] Chen, Y.-H., Lee, M.-J., Wang, I.-C., S. Du, Y.-F. C., Chen, Y.-C. & Yu, I. A. Coherent Optical Memory with High Storage Efficiency and Large Fractional Delay. *Physical Review Letters* **110**, 083601 (2013).
- [157] Eisaman, M. D., André, a., Massou, F., Fleischhauer, M., Zibrov, a. S. & Lukin, M. D. Electromagnetically induced transparency with tunable single-photon pulses. *Nature* **438**, 837–841 (2005).
- [158] Chaneliere, T., Matsukevich, D. N., Jenkins, S. D., Lan, S.-Y., T. A. B. Kennedy & Kuzmich, A. Storage and retrieval of single photons transmitted between remote quantum memories. *Nature* **438**, 833–836 (2005).
- [159] Appel, J., Figueroa, E., Korystov, D., Lobino, M. & Lvovsky, a. I. Quantum memory for squeezed light. *Physical Review Letters* **100**, 1–4 (2008).
- [160] Honda, K. *et al.* Storage and Retrieval of a Squeezed Vacuum. *Physical Review Letters* **100**, 093601 (2008).

- [161] Arikawa, M., Honda, K., Akamatsu, D., Nagatsuka, S., Akiba, K., Furusawa, A. & Kozuma, M. Quantum memory of a squeezed vacuum for arbitrary frequency sidebands. *Physical Review A* **81**, 021605 (2010).
- [162] Choi, K. S., Deng, H., Laurat, J. & Kimble, H. J. Mapping photonic entanglement into and out of a quantum memory. *Nature* **452**, 67–71 (2008).
- [163] Hétet, G. *et al.* Erratum: Squeezing and entanglement delay using slow light. *Physical Review A* **74**, 059902(E) (2006).
- [164] Hétet, G., Peng, A., Johnsson, M. T., Hope, J. J. & Lam, P. K. Characterization of electromagnetically-induced-transparency-based continuous-variable quantum memories. *Physical Review Letters* **77**, 12323 (2008).
- [165] Hsu, M. T. L., Hétet, G., Glöckl, O., Longdell, J. J., Buchler, B. C., H.-A. Bachor & Lam, P. K. Quantum Study of Information Delay in Electromagnetically Induced Transparency. *Physical Review Letters* **97**, 183601 (2006).
- [166] Figueroa, E., Vewinger, F., Appel, J. & Lvovsky, A. I. Decoherence of electromagnetically induced transparency in atomic vapor. *Optics Letters* **31**, 2625–2627 (2006).
- [167] Smekal, A. Zur Quantentheorie der Dispersion. *Naturwissenschaften* **11**, 873–875 (1923).
- [168] Raman, C. V. A new radiation. *Indian J. Phys.* **2**, 387–398 (1928).
- [169] Landsberg, G. & Mandelstam, L. Über die Lichtzerstreuung in Kristallen. *Zeitschrift für Physik* **50**, 769 (1928).
- [170] Nunn, J. *Quantum Memory in Atomic Ensembles*. Ph.D. thesis (2007).
- [171] Gorshkov, A. V., André, A., Lukin, M. D. & Sørensen, A. S. Photon storage in Λ -type optically dense atomic media. I. Cavity model. *Physical Review A* **76**, 033804 (2007).
- [172] Nunn, J., Walmsley, I. A., Raymer, M. G., Surmacz, K., Waldermann, F. C., Wang, Z. & Jaksch, D. Mapping broadband single-photon wave packets into an atomic memory. *Physical Review A* **75** (2007).
- [173] Kozhokin, a. E., Molmer, K. & Polzik, E. Quantum Memory for Light **62**, 1–5 (1999).
- [174] Reim, K. F. *et al.* Towards high-speed optical quantum memories. *Nature Photonics* **4**, 218–221 (2010).
- [175] Camacho, R. M., Vudyasetu, P. K. & Howell, J. C. Four-wave-mixing stopped light in hot atomic rubidium vapour. *Nature Photonics* **3**, 103–106 (2009).
- [176] Bustard, P. J., Lausten, R., England, D. G. & Sussman, B. J. Toward quantum processing in molecules: A THz-bandwidth coherent memory for light. *Physical Review Letters* **111**, 1–5 (2013).
- [177] England, D., Bustard, P., Nunn, J., Lausten, R. & Sussman, B. From photons to phonons and back: A thz optical memory in diamond. *Physical Review Letters* **111**, 243601 (2013).
- [178] England, D. G., Fisher, K. A. G., Jean-Philippe W. MacLean, P. J. B., Lausten, R., Resch, K. J. & Sussman, B. J. Storage and Retrieval of THz-Bandwidth Single Photons Using a Room-Temperature Diamond Quantum Memory. *Physical Review Letters* **114**, 053602 (2015).
- [179] Hahn, E. L. Spin Echoes. *Physical Review* **80**, 580–594 (1950).

- [180] Kurnit, I. D., Abella, N. & Hartmann, S. R. Observation of a photon echo. *Physical Review Letters* **13** (1964).
- [181] Ledingham, P. M., Naylor, W. R., Longdell, J. J., Beavan, S. E. & Sellars, M. J. Nonclassical photon streams using rephased amplified spontaneous emission. *Physical Review A* **81**, 012301 (2010).
- [182] Beavan, S. E., Hedges, M. P. & Sellars, M. J. Demonstration of Photon-Echo Rephasing of Spontaneous Emission. *Physical Review Letters* **109**, 093603 (2012).
- [183] Ledingham, P. M., Naylor, W. R. & Longdell, J. J. Experimental Realization of Light with Time-Separated Correlations by Rephasing Amplified Spontaneous Emission. *Physical Review Letters* **109**, 093602 (2012).
- [184] Ruggiero, J.-L., Gouet, L., Simon, C. & Chanelière, T. Why the two-pulse photon echo is not a good quantum memory protocol. *Physical Review A* **79**, 053851 (2009).
- [185] Tittel, W., Afzelius, M., Chanelière, T., Cone, R., Kroll, S., Moiseev, S. & Sellars, M. Photon-echo quantum memory in solid state systems. *Laser Photon. Rev.* **4**, 244–267 (2010).
- [186] Moiseev, S. A. & Kröll, S. Complete Reconstruction of the Quantum State of a Single-Photon Wave Packet Absorbed by a Doppler-Broadened Transition. *Physical Review Letters* **87**, 173601 (2001).
- [187] Moiseev, S. A., Tarasov, V. F. & Ham, B. S. Quantum memory photon echo-like techniques in solids. *J. Opt. B: Quantum Semiclassical Opt.* **5**, 497 (2003).
- [188] Alexander, A. L. *et al.* Photon Echoes Produced by Switching Electric Fields. *Physical Review Letters* **96**, 43602 (2006).
- [189] Hétet, G., Hosseini, M., Sparkes, B. M., D. Oblak, P. K. L. & Buchler, B. C. Photon echoes generated by reversing magnetic field gradients in a rubidium vapor. *Optics Letters* **33**, 2323 (2008).
- [190] Nilsson, M. & Kroll, S. Solid state quantum memory using complete absorption and re-emission of photons by tailored and externally controlled inhomogeneous absorption profiles. *Optics Communications* **247**, 393–403 (2005).
- [191] Ohlsson, N., Mohan, R. K. & Kroll, S. Quantum computer hardware based on rare-earth-ion-doped inorganic crystals. *Optics Communications* **201**, 71–77 (2002).
- [192] Alexander, A. L., Longdell, J. J., Sellars, M. J. & Manson, N. B. Coherent information storage with photon echoes produced by switching electric fields. *Journal of Luminescence* **127**, 94–97 (2007).
- [193] Lauritzen, B., Minar, J., de Riedmatten, H., Afzelius, M., Sangouard, N., Simon, C. & Gisin, N. Telecommunication-Wavelength Solid-State Memory at the Single Photon Level. *Physical Review Letters* **104**, 080502 (2010).
- [194] Hedges, M. P., Longdell, J. J., Li, Y. & Sellars, M. J. Efficient quantum memory for light. *Nature* **465**, 1052–1056 (2010).
- [195] Hétet, G., Longdell, J. J., Alexander, A. L., Lam, P. K. & Sellars, M. J. Electro-Optic Quantum Memory for Light Using Two-Level Atoms. *Physical Review Letters* **100**, 23601 (2008).
- [196] Hosseini, M., Sparkes, B. M., Campbell, G., Lam, P. K. & Buchler, B. C. High efficiency coherent optical memory with warm rubidium vapour. *Nature Communications* **2**, 174 (2011).

- [197] de Riedmatten, H., Afzelius, M., Staudt, M. U., Simon, C. & Gisin, N. A solidstate light-matter interface at the single-photon level. *Nature* **456**, 773–777 (2008).
- [198] Sabooni, M., Li, Q., Kröll, S. & Rippe, L. Efficient quantum memory using a weakly absorbing sample. *Physical Review Letters* **110**, 1–5 (2013).
- [199] Bonarota, M., Le Gouët, J. L. & Chanelière, T. Highly multimode storage in a crystal. *New Journal of Physics* **13** (2011).
- [200] Saglamyurek, E. *et al.* Broadband waveguide quantum memory for entangled photons. *Nature* **469**, 512–515 (2011).
- [201] Clausen, C., Usmani, I., Bussi eres, F., Sangouard, N., Afzelius, M., de Riedmatten, H. & Gisin, N. Quantum storage of photonic entanglement in a crystal. *Nature* **469**, 508–511 (2011).
- [202] Usmani, I., Clausen, C., Bussieres, F., Sangouard, N., Afzelius, M. & Gisin, N. Heralded quantum entanglement between two crystals. *Nature Photonics* **6**, 234–237 (2012).
- [203] Afzelius, M. & Simon, C. Impedance-matched cavity quantum memory. *Physical Review A* **82**, 1–4 (2010).
- [204] Timoney, N., Usmani, I., Jobez, P., Afzelius, M. & Gisin, N. Single-photon-level optical storage in a solid-state spin-wave memory. *Physical Review A - Atomic, Molecular, and Optical Physics* **88** (2013).
- [205] G ndogan, M., Mazzera, M., Ledingham, P. M., Cristiani, M. & de Riedmatten, H. Coherent storage of temporally multimode light using a spin-wave atomic frequency comb memory. *New Journal of Physics* **15**, 045012 (2013).
- [206] G ndogan, M., Ledingham, P. M., Kutluer, K., Mazzera, M. & de Riedmatten, H. Solid State Spin-Wave Quantum Memory for Time-Bin Qubits. *Physical Review Letters* **114**, 230501 (2015).
- [207] Jobez, P., Laplane, C., Timoney, N., Gisin, N., Ferrier, A., Goldner, P. & Afzelius, M. Coherent Spin Control at the Quantum Level in an Ensemble-Based Optical Memory. *Physical Review Letters* **114**, 230502 (2015).
- [208] Gorshkov, A. V., Andr e, A., Lukin, M. D. & S rensen, A. S. Photon storage in Λ -type optically dense atomic media. II. Free-space model. *Phys. Rev. A* **76**, 33805 (2007).
- [209] Alzetta, G., Gozzini, a., Moi, L. & Orriols, G. An experimental method for the observation of r.f. transitions and laser beat resonances in oriented Na vapour. *Il Nuovo Cimento B Series 11* **36**, 5–20 (1976).
- [210] Gray, H. R., Whitley, R. M. & Stroud, C. R. Coherent trapping of atomic populations. *Optics letters* **3**, 218–220 (1978).
- [211] Parkins, A. S. Quantum-state mapping between multilevel atoms and cavity light fields. *Phys. Rev. A* **51**, 1578–1596 (1995).
- [212] Hennrich, M., Legero, T., Kuhn, a. & Rempe, G. Vacuum-stimulated Raman scattering based on adiabatic passage in a high-finesse optical cavity. *Physical Review Letters* **85**, 4872–4875 (2000).
- [213] He, G.-Z., Kuhn, A., Schieman, S. & Bergmann, K. Population transfer by stimulated Raman scattering with delayed pulses and by the stimulated-emission pumping method: a comparative study. *Journal of the Optical Society of America B* **7**, 1960 (1990).

- [214] Loudon, R. Fermion and boson beamsplitter statistics. *Physical Review A* **58** (1998).
- [215] Butcher, P. N. & Cotter, D. *Elements of Nonlinear Optics* (Cambridge University Press, Cambridge, 1990).
- [216] Loudon, R. *The Quantum Theory of Light* (Oxford University Press, 2000).
- [217] Zeuthen, E., Grodecka-Grad, A. & Sørensen, A. S. Three-dimensional theory of quantum memories based on Lambda-type atomic ensembles 19 (2011).
- [218] Holstein, T. & Primakoff, H. No Title. *Physical Review A* **58**, 1098–1113 (1940).
- [219] Koodynski, J., Chwedenczuk, J. & Wasilewski, W. Eigenmode description of Raman scattering in atomic vapors in the presence of decoherence. *Physical Review A* **86**, 013818 (2012).
- [220] Wasilewski, W. & Raymer, M. G. Pairwise entanglement and readout of atomic-ensemble and optical wave-packet modes in traveling-wave Raman interactions. *Conference on Lasers and Electro-Optics and 2006 Quantum Electronics and Laser Science Conference, CLEO/QELS 2006* (2006).
- [221] Braunstein, S. L. Squeezing as an irreducible resource. *Physical Review A - Atomic, Molecular, and Optical Physics* **71**, 1–4 (2005).
- [222] Wasilewski, W., Lvovsky, a. I., Banaszek, K. & Radzewicz, C. Pulsed squeezed light: Simultaneous squeezing of multiple modes. *Physical Review A - Atomic, Molecular, and Optical Physics* **73**, 1–12 (2006).
- [223] Reim, K. F. *et al.* Multipulse addressing of a Raman quantum memory: Configurable beam splitting and efficient readout. *Physical Review Letters* **108**, 1–4 (2012).
- [224] Sprague, M. R. *et al.* Efficient optical pumping and high optical depth in a hollow-core photonic-crystal fibre for a broadband quantum memory. *New Journal of Physics* **15** (2013).
- [225] Sprague, M. R. *et al.* Broadband single-photon-level memory in a hollow-core photonic crystal fibre. *Nature Photonics* **8**, 287–291 (2014).
- [226] Reim, K. F., Michelberger, P., Lee, K. C., Nunn, J., Langford, N. K. & Walmsley, I. A. Single-photon-level quantum memory at room temperature. *Physical Review Letters* **107**, 053603 (2011).
- [227] Essen, L. & Parry, J. V. L. No TitleAn atomic standard of frequency and time interval: A caesium resonator. *Nature* **176**, 280–282 (1955).
- [228] Sansonetti, J. E. & Martin, W. C. *Handbook of Basic Atomic Spectroscopy Data* (National Institute of Standards and Technology, 2003).
- [229] Steck, D. A. Cesium D Line Data. *Report* (2010).
- [230] Gerginov, V., Derevianko, A. & Tanner, C. E. Observation of the Nuclear Magnetic Octupole Moment of 133 Cs. *Physical Review Letters* **91**, 072501 (2003).
- [231] Rosenberry, M. A., Reyes, J. P., Tupa, D. & Gay, T. J. Radiation trapping in rubidium optical pumping at low buffer-gas pressures. *Physical Review A* **75**, 023401 (2007).
- [232] Ricci, L. *et al.* A compact grating-stabilized diode laser system for atomic physics. *Optics Communications* **117**, 541–549 (1995).

- [233] Hawthorn, C. J., Weber, K. P. & Scholten, R. E. Littrow configuration tunable external cavity diode laser with fixed direction output beam. *Review of Scientific Instruments* **72**, 4477–4479 (2001).
- [234] Smith, D. A. The role of hyper ne pumping in multilevel systems exhibiting saturated absorption. *American Journal of Physics* **72** (2004).
- [235] MacAdam, K. B. A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for cs and rb. *American Journal of Physics* **60** (1992).
- [236] Spring, J. B. *Integrated Photonics for Quantum Enhanced Technologies*. Ph.D. thesis (2013).
- [237] Surmacz, K., J. Nunn, K. R., Lee, K. C., Lorenz, V. O., Sussman, B., A. Walmsley, I. & Jaksch, D. Efficient spatially resolved multimode quantum memory. *Physical Review A (Atomic, Molecular, and Optical Physics)* **78** (2008).
- [238] Young, L. *et al.* Precision lifetime measurements of Cs 6p 2P_{1/2} and 6p 2P_{3/2} levels by single-photon counting. *Physical Review A* **50**, 2174 (1994).
- [239] Phillips, N. B., Gorshkov, A. V. & Novikova, I. Slow light propagation and amplification via electromagnetically induced transparency and four-wave mixing in an optically dense atomic vapor. *Journal of Modern Optics* **59**, 1916 (2009).
- [240] Lukin, M., Matsko, A., Fleischhauer, M. & Scully, M. Quantum noise and correlations in resonantly enhanced wave mixing based on atomic coherence. *Physical Review Letters* **82**, 1847–1850 (1999).
- [241] Reim, K. F. *Broadband Optical Quantum Memory*. Ph.D. thesis (2011).
- [242] Lund, A. P. & Ralph, T. C. Nondeterministic gates for photonic single-rail quantum logic. *Physical Review A* **66**, 032307 (2002).
- [243] Berry, D. W., Lvovsky, A. I. & Sanders, B. C. Interconvertibility of single-rail optical qubits. *Optics Letters* **31**, 107–109 (2006).
- [244] Julsgaard, B. & Mølmer, K. Fidelity of Fock-state-encoded qubits subjected to continuous-variable Gaussian processes. *Physical Review A* **89**, 12333 (2014).
- [245] Politi, A., Cryan, M. J., Rarity, J. G., Yu, S. & O’Brien, J. Silica-on-silicon waveguide quantum circuits. *Science* **320**, 646 (2008).
- [246] Humphreys, P. C., Metcalf, B. J., Spring, J. B., Moore, M., Jin, X.-M., Barbieri, M. & Walmsley, W. S. K. I. A. Linear Optical Quantum Computing in a Single Spatial Mode. *Physical Review Letters* **111**, 150501 (2013).
- [247] Marcikic, I., de Riedmatten, H., Tittel, W., Scarani, V., Zbinden, H. & Gisin, N. Time-bin entangled qubits for quantum communication created by femtosecond pulses. *Physical Review A* **66**, 062308 (2002).
- [248] Huntington, E. & Ralph, T. Components for optical qubits encoded in sideband modes. *Physical Review A* **69**, 1–5 (2004).
- [249] O’Brien, J. L., Pryde, G. J., White, A. G., Ralph, T. C. & Branning, D. Demonstration of an all-optical quantum controlled-not gate. *Nature* **426**, 6964 (2003).
- [250] Zhou, Z. Q., Lin, W. B., Yang, M., Li, C. F. & Guo, G. C. Realization of reliable solid-state quantum memory for photonic polarization qubit. *Physical Review Letters* **108**, 1–13 (2012).

- [251] Foot, C. J. *Atomic Physics* (Oxford University Press, 2005), first edn.
- [252] Cho, Y.-W. & Kim, Y.-H. Atomic vapor quantum memory for a photonic polarization qubit. *Optics express* **18**, 25786–25793 (2010).
- [253] Langford, N. K. *Encoding, manipulating and measuring quantum information in optics*. Ph.D. thesis (2007).
- [254] Mohseni, M. Quantum process tomography: Resource analysis of different strategies. *Physical Review A* **77**, 032322 (2008).
- [255] Lobino, M., Korystov, D., Kupchak, C., Figueroa, E., Sanders, B. C. & Lvovsky, A. I. Complete characterization of quantum-optical processes. *Science* **322**, 563–566 (2008).
- [256] Schwinger, J. Unitary Operator Bases. *Proc. Natl. Acad. Sci. U.S.A.* **46**, 570 (1960).
- [257] Wootters, W. K. Optimal state-determination by mutually unbiased measurements. *Annals of Physics* **191**, 363 (1989).
- [258] Adamson, R. B. A. & Steinberg, A. M. Improving quantum state estimation with mutually unbiased bases. *Physical Review Letters* **105**, 030406 (2010).
- [259] de Burgh, M. D. & Langford, N. K. Choice of measurement sets in qubit tomography. *Physical Review A* **78**, 052122 (2008).
- [260] O’Brien, J., Pryde, G., Gilchrist, A., James, D., Langford, N., Ralph, T. & White, A. Quantum process tomography of a controlled-not gate. *Physical Review Letters* **93**, 080502 (2004).
- [261] Kim, Y.-S., Cho, Y.-W., Ra, Y.-S. & Kim, Y.-H. Reversing the weak quantum measurement for a photonic qubit. *Optics Express* **17**, 11978 (2009).
- [262] Bhandari, R. & Samuel, J. Observation of topological phase by use of a laser interferometer. *Physical Review Letters* **60**, 1211 (1988).
- [263] Kim, Y.-H., Kulik, S. P. & Shih, Y. Quantum teleportation of a polarization state with a complete bell state measurement. *Physical Review Letters* **86**, 1370 (2001).
- [264] Bimbard, E. *et al.* Homodyne tomography of a single photon retrieved on demand from a cavity-enhanced cold atom memory. *Physical Review Letters* **112** (2014).
- [265] Rieländer, D., Kutluer, K., Ledingham, P. M., Gündogan, M., Fekete, J., Mazzer, M. & De Riedmatten, H. Quantum storage of heralded single photons in a praseodymium-doped crystal. *Physical Review Letters* **112**, 1–11 (2014).
- [266] Bussièrès, F. *et al.* Quantum teleportation from a telecom-wavelength photon to a solid-state quantum memory. *Nature* **8**, 775–778 (2014).
- [267] Jensen, K. *et al.* Quantum memory for entangled two-mode squeezed states. *Nature Physics* **7**, 5 (2010).
- [268] Fekete, J., Rieländer, D., Cristiani, M. & de Riedmatten, H. Ultranarrow-band photon-pair source compatible with solid state quantum memories and telecommunication networks. *Physical Review Letters* **110** (2013).
- [269] Tanzilli, S., Riedmatten, H. D., Zbinden, H., Baldi, P., De Micheli, M., Ostrowsky, D. & Gisin, N. Highly efficient photon-pair source using periodically poled lithium niobate waveguide. *Electronics Letters* **37**, 26–28 (2001).

- [270] Karpinski, M., Radzewicz, C. & Banaszek, K. Dispersion-based control of modal characteristics for parametric down-conversion in a multimode waveguide. *Optics Letters* **37**, 878 (2012).
- [271] Jachura, M., Karpinski, M., Radzewicz, C. & Banaszek, K. High-visibility nonclassical interference of photon pairs generated in a multimode nonlinear waveguide. *Optics Express* **22**, 8624–8632 (2014).
- [272] Alcock, C. B., Itkin, V. P. & Horrigan, M. K. Vapor Pressure Equations for the Metallic Elements: 2982500 K. *Canadian Metallurgical Quarterly* **23** (1984).
- [273] Scully, M. O. & Zubairy, S. *Quantum Optics* (Cambridge University Press, 1997).
- [274] Hanbury Brown, R. & Twiss, R. Q. Interferometry of the intensity fluctuations in light. I. Basic theory: the correlation between photons in coherent beams of radiation. *Proceedings of the Royal Society A* **242**, 300–324 (1957).
- [275] Abe, S. & Matsuo, Y. Phase-space representation of completely positive quantum operation: Invertible subdynamics of two-mode squeezed states. *Physical Review E* **91**, 012126 (2015).
- [276] Christ, A., Laiho, K., Eckstein, A., Cassemiro, K. N. & Silberhorn, C. Probing multimode squeezing with correlation functions. *New Journal of Physics* **13** (2011).
- [277] Wu, C. B., Raymer, M. G., Wang, Y. Y. & Benabid, F. Quantum theory of phase correlations in optical frequency combs generated by stimulated Raman scattering. *Physical Review A* **82**, 053834 (2010).
- [278] Raymer, M. & Mostowski, J. Stimulated raman scattering: unified treatment of spontaneous initiation and spatial propagation. *Physical Review A* **24** (1981).
- [279] Goldschmidt, E. a. *et al.* Mode reconstruction of a light field by multiphoton statistics. *Physical Review A - Atomic, Molecular, and Optical Physics* **88**, 1–5 (2013).
- [280] Mandel, L. & E.Wolf. *Optical Coherence and Quantum Optics* (Cambridge University Press, 1995).
- [281] Lvovsky, A. I. Squeezed Light. *Photonics: Scientific Foundations, Technology and Applications* **1** (2014).
- [282] Fang, Y. & Jing, J. Quantum squeezing and entanglement from a two-mode phase-sensitive amplifier via fourwave mixing in rubidium vapor. *New Journal of Physics* **17**, 023027 (2015).
- [283] Bogoliubov, N. N. No Title. *J. Phys. (U.S.S.R.)* **11** (1947).
- [284] Dabrowski, M., Chrapkiewicz, R. & Wasilewski, W. Hamiltonian design in readout from room-temperature Raman atomic memory. *Optics Express* **22**, 26076–26091 (2014).
- [285] Kumar, P. & Kolobov, M. I. Degenerate four-wave mixing as a source for spatially-broadband squeezed light. *Optics Communications* **104**, 374–378 (1994).
- [286] Lee, K. C. *et al.* Macroscopic non-classical states and terahertz quantum processing in room-temperature diamond. *Nature Photonics* **4**, 41–45 (2012).
- [287] Vurgaftman, I. & Bashkansky, M. Suppressing four-wave mixing in warm-atomic-vapor quantum memory. *Physical Review A - Atomic, Molecular, and Optical Physics* **87**, 1–6 (2013).

- [288] Collett, M. J. & Gardiner, C. Squeezing of intracavity and traveling-wave light fields produced in parametric amplification. *Physical Review A* **30**, 1386 (1984).
- [289] Gardiner, C. W. & Savage, C. M. A multimode quantum theory of a degenerate parametric amplifier in a cavity. *Optics Communications* **50** (1984).
- [290] Yurke, B. No Title. *Physical Review A* **29** (1984).
- [291] Barnett, S. M., Gilson, C. R., Huttner, B. & Imoto, N. Field Commutation Relations in Optical Cavities. *Physical Review Letters* **77** (1996).
- [292] Walls, D. F. & Milburn, G. J. *Quantum Optics* (Springer, Berlin, 2008).
- [293] Raymer, M. G. & McKinstrie, C. J. Quantum input-output theory for optical cavities with arbitrary coupling strength: Application to two-photon wave-packet shaping. *Physical Review A - Atomic, Molecular, and Optical Physics* **88**, 1–14 (2013).
- [294] Viviescas, C. & Hackenbroich, G. Field quantization for open optical cavities. *Physical Review A* **67**, 013805 (2003).
- [295] Lang, R., Scully, M. O. & W. E. Lamb, J. Why is the Laser Line So Narrow? A Theory of Single-Quasimode Laser Operation. *Physical Review A* **7**, 1788 (1973).
- [296] Roch, J.-F., Vigner, K., Grelu, P., Sinatra, A., Poizat, J.-P. & Grangier, P. Quantum Nondemolition Measurements using Cold Trapped Atoms. *Physical Review Letters* **78**, 634 (1997).
- [297] Kuzmich, A., Bigelow, N. P. & Mandel, L. Atomic quantum non-demolition measurements and squeezing. *Europhysics Letters* **42**, 481 (1998).
- [298] Gorshkov, A. V., André, A., Fleischhauer, M., Sørensen, A. S. & Lukin, M. D. Universal Approach to Optimal Photon Storage in Atomic Media. *Physical Review Letters* **98**, 123601 (2007).
- [299] Schleich, W. P. & Walther, H. *Elements of Quantum Information* (John Wiley & Sons, Weinheim, 2007).
- [300] Eckstein, A., Brecht, B. & Silberhorn, C. A quantum pulse gate based on spectrally engineered sum frequency generation. *Optics Express* **19**, 13770–13778 (2011).
- [301] Brecht, B., Eckstein, A., Ricken, R., Quiring, V., Suche, H., Sansoni, L. & Silberhorn, C. Demonstration of coherent time-frequency Schmidt mode selection using dispersion-engineered frequency conversion. *Physical Review A* **90**, 030302 (2014).
- [302] Reddy, D. V., Raymer, M. G. & McKinstrie, C. J. Efficient sorting of quantum-optical wave packets by temporal-mode interferometry. *Optics Letters* **39** (2014).
- [303] Christ, A., Brecht, B., Maurer, W. & Silberhorn, C. Theory of quantum frequency conversion and type-II parametric down-conversion in the high-gain regime. *New Journal of Physics* **15**, 053038 (2013).
- [304] Saleh, B. E. A. & Teich, M. C. *Fundamentals of Photonics* (John Wiley & Sons, 1991).
- [305] Ghosh, G. *Handbook of Thermo-Optic Coefficients of Optical Materials with Applications* (Academic Press, Waltham, 1998), first edn.

-
- [306] Hänsch, T. W. & Couillaud, B. Laser Frequency Stabilization by Polarization Spectroscopy of a Reflecting Reference Cavity. *Optics Communications* **35** (1980).
- [307] White, A. D. No Title. *IEE J. Quant. Electron.* **8**, 349 (1965).
- [308] Collett, E. *Field Guide to Polarization* (SPIE, 2005).
- [309] Harvey, M. D. & White, A. G. Frequency locking by analysis of orthogonal modes. *Optics Communications* **221**, 163–171 (2003).
- [310] Thorlabs. HRS015 Stabilized Red HeNe Laser User Guide. Tech. Rep. (2013).
- [311] Siddons, P., Adams, C. S. & Hughes, I. G. Off-resonance absorption and dispersion in vapours of hot alkali-metal atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics* **42**, 175004 (2009).
- [312] Siddons, P., Adams, C. S., Ge, C. & Hughes, I. G. Absolute absorption on rubidium D lines: comparison between theory and experiment. *Journal of Physics B: Atomic, Molecular and Optical Physics* **41**, 155004 (2008).
- [313] Faddeeva, V. N. & Terentev, N. M. *Tables of Values of the Error Function with a Complex Argument (Engl. Transl.)* (Pergamon, New York, 1961).
- [314] Zaghoul, M. R. Comment on a fast method of modeling spectral line. *J. Quant. Spectrosc. Radiat. Transfer* **109**, 2895 (2008).
- [315] Nunn, J. Personal communication (2015).
- [316] Hong, C., Ou, Z. & Mandel, L. Measurement of Subpicosecond Time Intervals between Two Photons by Interference. *Physical Review Letters* **59**, 2044–2046 (1987).
- [317] Lewis-Swan, R. J. & Kheruntsyan, K. V. Proposal for demonstrating the HongOuMandel effect with matter waves. *Nature Communications* **5** (2014).
- [318] Fattal, D. *Single Photons for Quantum Information Processing*. Ph.D. thesis (2010).
- [319] Ekert, A. K. & Knight, P. L. Correlations and Squeezing of Two-Mode Oscillations. *American Journal of Physics* **57**, 692–697 (1989).
- [320] Siegman, A. E. *Lasers* (University Science Books, 1986).

Appendix A

Second-order autocorrelation function

The photon number statistics of different states of classical and non-classical light can be distinguished by measuring the second-order intensity autocorrelation function $g^{(2)}$, as described in Section 5.3.3. We can define the normalised $g^{(2)}$ in terms of normally ordered products of photon creation and annihilation operators so that - at zero time delay -

$$g^{(2)}(\tau = 0) = \frac{\langle \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} \rangle}{\langle \hat{a}^\dagger \hat{a} \rangle^2}. \quad (\text{A.1})$$

This expression can be rewritten in terms of the mean $\langle n \rangle$ and variance V_n of the photon number distribution^[318], given by

$$\begin{aligned} V_n &= \langle \Delta n^2 \rangle = \langle n^2 \rangle - \langle n \rangle^2 \\ &= \langle (\hat{a}^\dagger \hat{a})^2 \rangle - \langle \hat{a}^\dagger \hat{a} \rangle^2 \\ &= \langle \hat{a}^\dagger (1 + \hat{a}^\dagger \hat{a}) \hat{a} \rangle - \langle \hat{a}^\dagger \hat{a} \rangle^2 \\ &= \langle \hat{a}^{\dagger 2} \hat{a}^2 \rangle + \langle \hat{a}^\dagger \hat{a} \rangle - \langle \hat{a}^\dagger \hat{a} \rangle^2. \end{aligned} \quad (\text{A.2})$$

Re-arranging this equation and normalising by $\langle n \rangle$ gives

$$\frac{V_n - \langle n \rangle}{\langle n \rangle} = \frac{\langle \hat{a}^{\dagger 2} \hat{a}^2 \rangle}{\langle \hat{a}^{\dagger} \hat{a} \rangle} - \langle \hat{a}^{\dagger} \hat{a} \rangle = \langle n \rangle \left(g^{(2)}(0) - 1 \right). \quad (\text{A.3})$$

This quantity is known as *Mandel's Q parameter*^[319].

A.1 Fock states

The photon number states $|n\rangle$ are simultaneous eigenstates of the number operator $\hat{n} = \hat{a}^{\dagger} \hat{a}$ and the Hamiltonian for a one-dimensional harmonic oscillator, i.e.

$$\begin{aligned} \hat{n} |n\rangle &= \hat{a}^{\dagger} \hat{a} |n\rangle = n |n\rangle \\ \hat{H} |n\rangle &= \hbar\omega \left(\hat{a}^{\dagger} \hat{a} + \frac{1}{2} \right) |n\rangle = \left(n + \frac{1}{2} \right) \hbar\omega |n\rangle = E_n |n\rangle. \end{aligned} \quad (\text{A.4})$$

By applying the creation operator \hat{a}^{\dagger} to the second of these equations and making use of the commutator $[\hat{a}, \hat{a}^{\dagger}] = 1$, we find that

$$\hbar\omega \left(\hat{a}^{\dagger} \hat{a}^{\dagger} \hat{a} + \frac{1}{2} \hat{a}^{\dagger} \right) |n\rangle = \hbar\omega \left(\hat{a}^{\dagger} \hat{a} \hat{a}^{\dagger} - \hat{a}^{\dagger} + \frac{1}{2} \hat{a}^{\dagger} \right) |n\rangle = E_n \hat{a}^{\dagger} |n\rangle, \quad (\text{A.5})$$

which can be rearranged to give

$$\hbar\omega \left(\hat{a}^{\dagger} \hat{a} + \frac{1}{2} \right) \hat{a}^{\dagger} |n\rangle = (E_n + \hbar\omega) \hat{a}^{\dagger} |n\rangle. \quad (\text{A.6})$$

$\hat{a}^{\dagger} |n\rangle$ is therefore an eigenstate of \hat{H} with eigenvalue $E_n + \hbar\omega$ so that we can write $C_{n+1} |n+1\rangle = \hat{a}^{\dagger} |n\rangle$ and likewise, $C_n |n-1\rangle = \hat{a} |n\rangle$. The normalisation constants are

given by $C_{n+1} = \sqrt{n+1}$ and $C_n = \sqrt{n}$. Hence,

$$\begin{aligned}\hat{a}^\dagger |n\rangle &= (n+1)^{1/2} |n+1\rangle \\ \hat{a} |n\rangle &= n^{1/2} |n-1\rangle.\end{aligned}\tag{A.7}$$

Using equations (A.7), the $g^{(2)}$ for a single-mode state $|n\rangle$ is given by

$$g^{(2)} = \frac{\langle n | \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} | n \rangle}{\langle n | \hat{a}^\dagger \hat{a} | n \rangle^2} = \frac{n(n-1)}{n^2} = 1 - \frac{1}{n} < 1.\tag{A.8}$$

This inequality corresponds to sub-Poissonian statistics for which the variance of the photon number distribution is less than its mean (see equation (A.3)), thus violating the Cauchy-Schwartz inequality for classical distributions. A signature of this is the anti-bunching effect observed in Hanbury Brown-Twiss-type measurements, revealing the anti-correlation of non-classical photon states. For a single photon state $n = 1$, we therefore find $g^{(2)} = 0$.

A.2 Coherent states

A coherent state $|\alpha\rangle$ is unchanged by the detection (or annihilation) of a photon and is thus defined according to $\hat{a} |\alpha\rangle = |\alpha\rangle$. By introducing the unitary displacement operator $D(\alpha) = \exp(\alpha \hat{a}^\dagger - \alpha^* \hat{a})$, we can re-express the coherent state in terms of the vacuum state $|0\rangle$ as $|\alpha\rangle = D(\alpha) |0\rangle$. It can be shown^[273] that

$$\begin{aligned}D^\dagger(\alpha) \hat{a} D(\alpha) &= \hat{a} + \alpha \\ D^\dagger(\alpha) \hat{a}^\dagger D(\alpha) &= \hat{a}^\dagger + \alpha^*.\end{aligned}\tag{A.9}$$

These relations demonstrate the action of the displacement operator which can be visualised as a displacement of a localised state by an amount α in phase space. The coherent state can therefore be understood as a displaced form of the ground state $|0\rangle$ of the harmonic oscillator^[273].

Using the displacement properties (A.9), we can now determine the mean and variance of the coherent state photon number distribution:

$$\begin{aligned}\langle n \rangle &= \langle 0 | D^\dagger \hat{a}^\dagger \hat{a} D | 0 \rangle = \langle 0 | (\hat{a}^\dagger D^\dagger + \alpha^* D^\dagger) \hat{a} D | 0 \rangle \\ &= \langle 0 | \hat{a}^\dagger (\hat{a} + \alpha) + \alpha^* (\hat{a} + \alpha) | 0 \rangle = \langle 0 | \alpha^* \alpha | 0 \rangle \\ &= |\alpha|^2\end{aligned}\tag{A.10}$$

Similarly, we find for the variance V_n , $(\Delta n)^2 = \langle n^2 \rangle - \langle n \rangle^2 = |\alpha|^2$.

As expected for a photon number distribution governed by Poissonian statistics, the variance V_n is equal to the mean $\langle n \rangle$. Then, according to equation (A.3), $g^{(2)} = 1$ for coherent states^[273].

A.3 Thermal states

As discussed in Section 5.4, the photon number statistics of the measured FWM noise contribution are expected to resemble thermal statistics as a result of tracing over one of the modes of a two-mode squeezed state. Thermal statistics are characterised by an exponential probability distribution which we can examine by considering a single-mode thermal state in equilibrium with a blackbody at temperature T . The probability p_n of occupying a state

with energy E_n follows a Boltzmann distribution, i.e.

$$p_n = \frac{\exp(-\beta E_n)}{\sum_n \exp(-\beta E_n)}, \quad (\text{A.11})$$

where $E_n = (n + \frac{1}{2}) \hbar\omega$ is the n^{th} energy eigenstate of the harmonic oscillator and $\beta = 1/k_B T$. The density of states ρ_{th} can be expressed as

$$\rho_{\text{th}} = \frac{\exp(-\beta E_n)}{\text{tr}(\exp(-\beta E_n))} = \frac{1}{Z} \exp(-\beta E_n) = \sum_{n=0}^{\infty} p_n |n\rangle \langle n|. \quad (\text{A.12})$$

Here, the partition function Z , or sum over states¹, normalises the probability function and is equal to

$$\begin{aligned} Z &= \text{tr}(\exp(-\beta E_n)) = \sum_{n=0}^{\infty} \langle n | \exp(-\beta \hat{H}) | n \rangle \\ &= \sum_{n=0}^{\infty} \exp(-\beta E_n) = \frac{\exp(-\beta \frac{\hbar\omega}{2})}{1 - \exp(-\beta \hbar\omega)} \end{aligned} \quad (\text{A.13})$$

The mean photon number can be found by evaluating

$$\begin{aligned} \langle n \rangle &= \text{tr}(n \rho_{\text{th}}) = \sum_{n=0}^{\infty} \langle n | n \rho_{\text{th}} | n \rangle \\ &= \sum_{n=0}^{\infty} n p_n = \exp\left(-\beta \frac{\hbar\omega}{2}\right) \frac{1}{Z} \sum_{n=0}^{\infty} n \exp(-\beta \hbar\omega n). \end{aligned} \quad (\text{A.14})$$

Using the identity $\sum_{n=0}^{\infty} n \exp(-nx) = -\frac{d}{dx} \sum_{n=0}^{\infty} \exp(-nx)$, we find that the mean number of photons obeys the Planck distribution,

$$\langle n \rangle = \frac{1}{\exp(\hbar\omega\beta) - 1}, \quad (\text{A.15})$$

as expected for a blackbody in thermal equilibrium.

¹In German: *Zustandssumme*.

By re-arranging the above equation to $\exp(-\hbar\omega\beta) = \langle n \rangle / (\langle n \rangle + 1)$, the photon number distribution can be expressed in terms of the mean photon number as

$$p_n = \frac{\langle n \rangle^n}{(1 + \langle n \rangle)^{n+1}}. \quad (\text{A.16})$$

The expectation value $\langle n^2 \rangle$ can be determined as follows:

$$\begin{aligned} \langle n^2 \rangle &= \text{tr}(n^2 \rho_{\text{th}}) = \sum_{n=0}^{\infty} \langle n | n^2 \rho_{\text{th}} | n \rangle \\ &= \sum_{n=0}^{\infty} n^2 p_n = \exp\left(-\beta \frac{\hbar\omega}{2}\right) \frac{1}{Z} \sum_{n=0}^{\infty} n^2 \exp(-\beta \hbar\omega n). \end{aligned} \quad (\text{A.17})$$

We can evaluate this sum by noting that

$$\begin{aligned} \sum_{n=0}^{\infty} n^2 e^{-nx} &= \frac{d}{dx} \left(\frac{d}{dx} \sum_{n=0}^{\infty} e^{-nx} \right) = \frac{d}{dx} \left(\frac{-e^{-x}}{(1 - e^{-x})^2} \right) \\ &= \frac{e^{-x} + e^{-2x}}{(1 - e^{-x})^3}. \end{aligned} \quad (\text{A.18})$$

Hence,

$$\begin{aligned} \langle n^2 \rangle &= (1 - e^{-x}) \frac{e^{-x} + e^{-2x}}{(1 - e^{-x})^3} \\ &= \frac{e^{-x}}{1 - e^{-x}} + \frac{2e^{-2x}}{(1 - e^{-x})^2} \\ &= \langle n \rangle + 2\langle n \rangle^2 \end{aligned} \quad (\text{A.19})$$

The variance is then given by $\langle (\Delta n)^2 \rangle = \langle n^2 \rangle - \langle n \rangle^2 = \langle n \rangle + \langle n \rangle^2$ and according to equation (A.3), $g^{(2)} = 2$ for a single-mode thermal state.

Appendix **B**

The cavity stability condition

B.1 Ray transfer matrix

In the formalism of ray transfer matrix analysis, the ABCD matrix characterises the change in the position and direction of an incident beam as it propagates between two reference planes, known as *input* and *output* planes. These planes are defined to be perpendicular to the optical axis of the system which – in the case of lenses and curved mirrors – is chosen to pass through the centre of curvature. The input ray can be represented by the vector $\mathbf{r}_1 = (r_1 \ \theta_1)^T$, where r_1 is the distance from the optical axis, and θ_1 is the angle between the optical axis and the direction of propagation.

After passing through the optical system, the outgoing ray at the output plane can be found from the transformation

$$\begin{pmatrix} r_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} r_1 \\ \theta_1 \end{pmatrix}. \quad (\text{B.1})$$

As an aside, the coordinates of a ray reflected from a curved mirror are related to the

input coordinates according to

$$\begin{aligned} r_2 &= r_1 \\ \theta_2 &= -(1/f)r_1 + \theta_1, \end{aligned} \tag{B.2}$$

where the focal length is given by $f = 2/R$ and R is the radius of curvature. The matrix coefficients are therefore equal to $A = 1$, $B = 0$, $C = -2/R$ and $D = 1$.

B.2 Stability analysis

The stability of an optical cavity, which periodically re-focuses the propagating beam, can be determined from the properties of the cavity's ABCD matrix^[320]. After n passes of the cavity, the ray vector \mathbf{r}_n is given by

$$\mathbf{r}_n = M^n \mathbf{r}_0, \tag{B.3}$$

where M is the matrix for one pass and \mathbf{r}_0 is the input ray. The stability condition can be derived by finding the eigenvalues λ of M , i.e.

$$M\mathbf{r} = \lambda\mathbf{r}, \tag{B.4}$$

or equivalently,

$$(M - \lambda\mathbf{I})\mathbf{r} = 0, \tag{B.5}$$

where \mathbf{I} is the 2×2 identity matrix.

For non-trivial solutions of \mathbf{r} , we set the determinant to zero, giving the characteristic

equation

$$\det(M - \lambda \mathbf{I}) = \lambda^2 - \text{tr}(M)\lambda + \det(M) = 0. \quad (\text{B.6})$$

Using the property $\det(M) = AD - BC = 1$, we find

$$\lambda^2 - (A + D)\lambda + 1 = 0, \quad (\text{B.7})$$

with eigenvalue solutions

$$\lambda_{\pm} = m \pm \sqrt{m^2 - 1}. \quad (\text{B.8})$$

Here, we have defined the parameter $m = \frac{1}{2}\text{tr}(M) = (A + D)/2$.

Having determined the eigenvalues using equation (B.8), the input ray can be expressed as a linear superposition of the orthogonal eigenvectors \mathbf{r}_{\pm} of M . These *eigenrays* can be found from $M\mathbf{r}_{+} = \lambda_{+}\mathbf{r}_{+}$ and $M\mathbf{r}_{-} = \lambda_{-}\mathbf{r}_{-}$, respectively. Inserting $\mathbf{r}_0 = c_{+}\mathbf{r}_{+} + c_{-}\mathbf{r}_{-}$ into equation (B.3) gives

$$\begin{aligned} \mathbf{r}_n &= M^n \times (c_{+}\mathbf{r}_{+} + c_{-}\mathbf{r}_{-}) \\ &= \lambda_{+}^n c_{+}\mathbf{r}_{+} + \lambda_{-}^n c_{-}\mathbf{r}_{-}, \end{aligned} \quad (\text{B.9})$$

where c_{\pm} are constants.

Now, if $|m| \leq 1$, we can set $m = \cos \theta$ so that the eigenvalues are given by the complex numbers

$$\lambda_{\pm} = \cos \theta \pm i \sin \theta = e^{\pm i\theta}. \quad (\text{B.10})$$

Hence, after the n^{th} round-trip of the cavity,

$$\begin{aligned} \mathbf{r}_n &= e^{in\theta} c_{+}\mathbf{r}_{+} + e^{-in\theta} c_{-}\mathbf{r}_{-} \\ &= \mathbf{r}_0 \cos(\theta n) + \mathbf{s}_0 \sin(\theta n), \end{aligned} \quad (\text{B.11})$$

where we have defined the “slope vector”^[320] $\mathbf{s}_0 = i(c_+\mathbf{r}_+ - c_-\mathbf{r}_-)$.

Equation (B.11) shows that the ray vector \mathbf{r}_n oscillates with a maximum displacement bounded by the sum of the initial vectors, $\mathbf{r}_0 + \mathbf{s}_0$. The condition $|m| \leq 1$ therefore represents a *stable* cavity. By writing this condition in terms of the ABCD matrix, we arrive at the stability criterion (6.54) used in Section 6.9.1

$$-1 \leq \frac{1}{2}\text{tr}(M) \leq +1. \quad (\text{B.12})$$

In the opposite case ($|m| > 1$), the eigenvalues λ_+, λ_- are real and positive quantities. The displacement r_n will therefore diverge, as the ray propagates around the cavity. This demonstrates that the condition $|m| > 1$ corresponds to an *unstable* cavity.