Imperial College London

Efficient, Low Noise, Mode-Selective Quantum Memory

Sarah Ellyn Thomas

Department of Physics
CDT in Controlled Quantum Dynamics

Supervised by
Joshua Nunn, Ian A. Walmsley & Myungshik S. Kim

Submitted for the degree of Doctor of Philosophy in Physics

June 2019

Declarations

I declare that the content of this thesis is my own work, and has not been submitted for any other degree. The results presented here have been obtained by myself, or in collaboration with my colleagues whom I acknowledge where appropriate. I certify that, to the best of my knowledge, I have referenced the work of others where appropriate.

Oxford, June 2019

Sarah Ellyn Thomas

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-NonCommercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Photonic quantum information processing is a key element for scalable quantum technologies, and has applications in secure long-distance quantum communication and connecting nodes of a quantum computation network. However, logical photon-photon gates and state-of-the-art single photon sources rely on probabilistic processes. Quantum memories are devices that enable storage and on-demand recall of quantum states of light, and have been highlighted as a vital component in photonic networks to overcome the scaling problem by synchronising probabilistic processes.

The Raman memory has a large storage bandwidth and high synchronising capacity, and is an ideal candidate for local synchronisation. However, previous demonstrations of the Raman memory suffer from four-wave mixing noise, which prohibits quantum level operation. In this thesis I investigate methods to increase the signal to noise ratio in the Raman memory. I investigate increasing the light-matter coupling strength to boost the memory efficiency, and then explore two different methods to suppress four-wave mixing noise. I demonstrate that operating the Raman memory in a cavity is successful in reducing four-wave mixing, but it is technically challenging to maintain a high memory efficiency. I investigate a new method of noise suppression by introducing an absorption feature at the frequency of the unwanted noise field. This technically simple method is successful in reducing the noise by an order of magnitude, and will be applicable to many quantum memory protocols.

In the final section of this thesis I explore the temporal mode properties of the Raman memory. I demonstrate that the Raman memory is single mode and can be used to separate and manipulate temporal modes of light. This positions the Raman memory as a key device for enabling high-dimensional photonic quantum information processing, and enhancing light-matter interactions.

These results pave the way towards an efficient, low-noise, mode-selective quantum memory.

Acknowledgements

I have been incredibly lucky to be surrounded by such wonderful support throughout my PhD. Thank you to everyone who has helped me get to this point.

Firstly I would like to thank Josh for your guidance and support over these past four years. Thank you for introducing me to the world of quantum memories with wisdom, humour, whiteboard scribbles, endless lambda systems and many cups of tea! And also thanks to Ian for overseeing it all with incredible insight, patience, and the knack of knowing when pints are needed. You have built up a phenomenal team in Ultrafast, and I look forward to seeing what the ImProv can do – stand-up or otherwise.

I've been extremely lucky to work alongside the post-doc dream team, and honestly could not have asked for a better line-up. Patrick, Dylan and Benni: you have been awesome, thank you. Firstly a huge thanks to Dylan, for showing me the ropes in the lab and for many in-depth discussions about physics, life and everything in between, all with incredibly infectious energy and enthusiasm. Thanks so much for all your support both in and out of the lab. Thank you to Benni for introducing me to "quantum rainbows" with heaps of scientific integrity, knowledge, mad lab skills and wonderful, dry humour. And finally thanks to Patrick for everything – it's been magical. You're an awesome scientist, a great leader of the meme team, and the true Quantum Boss.

It has been an enormous pleasure to share Memories Corner with two of the best students and crossword-solvers around. A huge thanks to Joe for everything you've taught me, from cavities to temporal modes, and from python to pints. And Tom you've been an incredible addition to the team, bringing the magic and the memes. Thank you for all your help and support, and I can't wait to see where the experiment goes next. Thank you both for making my PhD so enjoyable and for all the TEAPARTI tea parties.

Thank you to the rest of the memories team, past and present, and in particular Jonas, Oscar, Shaobo and Kris. It has been amazing to work with all of you, and I have thoroughly enjoyed being part of the meme team. Good luck with all the future experiments! Thanks also to the rest of the ultrafast group for being such a wonderful, if somewhat nutty, community in which to do my PhD. Thanks for all the pubfests, tea breaks, quizzes, punting, frisbee games and spring break, and thanks in particular to the Famous Five and to Joelle for keeping me sane through it all.

Outside of work I have had an incredible support network. Firstly, an enormous thank you to my housies Liz, Katie, Nonie and Lars, who have truly made my time in Oxford. Thank you for a wonderful four years filled with baking, brewing, puzzling, hiking, gardening, knitting, gossiping, puppy-sitting and adventuring – cheers baes!

A huge thanks to the Cambridge–Oxford bunch, and in particular Anita, for the mountains of food, cups of tea and the (gradually more successful) pub quizzes! Thank you to Felicity, Grishma and Katie for years of wonderful friendship, encouragement, and sharing the same terrible taste in TV shows. And thank you to Jamie whose phenomenal support and encouragement means the world.

Lastly, but most importantly, thank you to my family. Thank you Oma for the perfect respite from PhD work, with endless tea, smiles and stories. Thank you Owen and Jo for being the best sibs and friends, and for being my biggest fans through every up and down. And finally an enormous thank you to my parents Karen and Hywel – your tireless love, support and generosity has got me here, and I'm more grateful than I could ever say.

Author Publications

Articles

Related Publications

"Raman Quantum Memory with Built-in Suppression of Four-wave Mixing Noise"
S. E. Thomas, T. M. Hird, J. H. D. Munns, B. Brecht, D. J. Saunders, J. Nunn, I. A. Walmsley, & P. M. Ledingham
Submitted, arXiv 1905.00042

"Temporal–Mode Selectivity With a Raman Memory"
J. H. D. Munns, S. E. Thomas, J. Sperling, P. M. Ledingham, D. J. Saunders, J. Nunn, B. Brecht & I. A. Walmsley
In preparation

"Theory of noise suppressions in Λ-type quantum memories by means of a cavity"
J. Nunn, J. H. D. Munns, S. E. Thomas, K. T. Kaczmarek, C. Qiu, A. Feizpour, E. Poem, B. Brecht, D. J. Saunders, P. M. Ledingham, D. V. Reddy, M. G. Raymer & I. A. Walmsley
Physical Review A **96**, 012338 (2017)

"High efficiency Raman memory by suppressing radiation trapping"
S. E. Thomas, J. H. D. Munns, K. T. Kaczmarek, C. Qiu, B. Brecht, A. Feizpour, P. M. Ledingham, I. A. Walmsley, J. Nunn & D. J. Saunders
New Journal of Physics **19**, 063034 (2017)

Other

"Optimal Coherent Filtering for Single Noisy Photons"
S. Gao, O. Lazo-Arjona, B. Brecht, K. T. Kaczmarek, P. M. Ledingham, S. E. Thomas, J. Nunn, D. J. Saunders & I. A. Walmsley
Submitted, arXiv 1902.07720

"High-speed noise-free optical quantum memory"
K. T. Kaczmarek, P. M. Ledingham, B.Brecht, S. E. Thomas, G. S. Thekkadath, O. Lazo-Arjona, J. H. D. Munns, E. Poem, A. Feizpour, D. J. Saunders, J. Nunn, & I. A. Walmsley
Physical Review A **97**, 042316 (2018)

Conference Papers

Oral Presentations

"High-dimensional Temporal Mode Manipulation using Quantum Memories"
S. E. Thomas, J. H. D. Munns, M. Hynes, T. M. Hird, P. M. Ledingham, D. J. Saunders, J. Nunn, B. Brecht & I. A. Walmsley
Quantum Controlled Ultrafast Multimode Entanglement and Measurement (QCUMbER), Oxford (2018)

"A Noise-free Optical Quantum Memory at Room Temperature"
S. E. Thomas, K. T. Kaczmarek, P. M. Ledingham, B. Brecht, G. S. Thekkadath, O. Lazo-Arjona, J. H. D. Munns, E. Poem, A. Feizpour, D. J. Saunders, J. Nunn, & I. A. Walmsley.
Conference for Quantum Information and Quantum Control (CQIQC-VII), Toronto (2017)

"Storage of Nanosecond Heralded Single Photons in a Cavity-Enhanced Room-Temperature Broadband Quantum Memory"
S. E. Thomas, D. J. Saunders, J. H. D. Munns, T. F. M. Champion, C. Qiu, B. Brecht, K. T. Kaczmarek, E. Poem, P. M. Ledingham, I. A. Walmsley, & J. Nunn.
Young Quantum Information Scientists (YQIS), Barcelona (2016)

Poster Presentations

"High-dimensional Temporal Mode Manipulation using Quantum Memories"
S. E. Thomas, J. H. D. Munns, M. Hynes, T. M. Hird, P. M. Ledingham, D. J. Saunders, J. Nunn, B. Brecht & I. A. Walmsley
International Conference on Atomic Physics (ICAP), Barcelona (2018)

"High-dimensional Temporal Mode Manipulation using Quantum Memories"
S. E. Thomas, J. H. D. Munns, K. T. Kaczmarek, P. M. Ledingham, D. J. Saunders, J. Nunn, B. Brecht & I. A. Walmsley
Quantum Communication, Measurement and Computing (QCMC), Louisiana (2018)

"Noise-free Broadband Quantum Memories for Temporal Multiplexing"
S. E. Thomas, K. T. Kaczmarek, P. M. Ledingham, B. Brecht, A. Feizpour, G. S. Thekkadath, J. H. D. Munns, D. J. Saunders, I. A. Walmsley & J. Nunn
Quantum Optics to Quantum Technology (QOtoQT), London (2017)

"High Efficiency Room-Temperature Raman Memory using Quenching"
S. E. Thomas, P. M. Ledingham, B. Brecht, J. H. D. Munns, C. Qiu, A. Feizpour, I. A. Walmsley, J. Nunn, & D. J. Saunders
Quantum Communication, Measurement and Computing (QCMC), Singapore (2016)

Contents

Declarations . . . 3
Abstract . . . 5
Acknowledgements . . . 6
Author Publications . . . 8
List of Figures . . . 12
List of Tables . . . 14
Abbreviations, Symbols and Constants . . . 15

1 Introduction **19**
1.1 Quantum-Enhanced Technology . . . 19
1.2 Architectures for Quantum Technologies . . . 21
1.3 Quantum Memories . . . 26

2 The Raman Quantum Memory **36**
2.1 Raman Memory Protocol . . . 36
2.2 Theory . . . 38
2.2.1 Optical Bloch Equations . . . 38
2.2.2 Modal Properties of the Raman Memory . . . 45
2.2.3 Beam Splitter Memory Interaction . . . 47
2.2.4 Four-wave mixing Noise . . . 49
2.3 Thesis Outline . . . 50
2.3.1 Increasing the Memory Efficiency . . . 50
2.3.2 Suppressing Four-Wave Mixing Noise . . . 51
2.3.3 Exploring Temporal Mode Properties . . . 53
2.4 Summary . . . 53

3 High Gain Raman Memory **54**
3.1 Experimental Set-up of the Raman Memory . . . 54
3.1.1 Storage Medium . . . 55
3.1.2 Signal and Control Field Preparation . . . 57
3.1.3 Filtering and Detection . . . 61
3.2 Optical Pumping . . . 63
3.3 High Temperature Raman Memory . . . 65
3.4 Summary . . . 69

4 Storage of Single Photons in a Cavity-Enhanced Raman Memory **70**
4.1 Theory of a Cavity-Enhanced Raman Memory . . . 70
4.2 Experimental Implementation of a Cavity Memory . . . 76
4.3 Second-Generation Cavity Memory . . . 79

4.4 Single Photon Source . . . 85
4.4.1 SPDC Photon Source . . . 85
4.4.2 Photon Source Filtering . . . 86
4.4.3 Heralding Efficiency . . . 87
4.4.4 Feed Forward . . . 88
4.5 Interfacing the Cavity-Enhanced Raman Memory with a Single Photon Source . 89
4.6 Summary . . . 90

5 Built-in Noise Suppression in a Raman Quantum Memory 92
5.1 Built-in Noise Suppression . . . 92
5.2 Theory and Simulations . . . 94
5.3 Experimental Results . . . 99
5.3.1 Efficiency and Storage Time . . . 100
5.3.2 Noise Characterisation . . . 101
5.4 Correlation Functions . . . 103
5.5 Summary . . . 108

6 Temporal Mode Properties of the Raman Memory 110
6.1 Temporal Modes . . . 110
6.1.1 Temporal Modes as a Basis for QIP . . . 111
6.1.2 Enhancing Light-Matter Interactions . . . 112
6.1.3 Devices for Manipulating Temporal Modes . . . 112
6.2 Temporal Mode Raman Memory Set-up . . . 114
6.2.1 Pulse Carver . . . 114
6.2.2 Tapered Amplifier . . . 115
6.2.3 Temporal Mode Raman Memory . . . 117
6.3 Mode Selectivity . . . 118
6.4 Mode Manipulation . . . 123
6.4.1 High Dimensional Temporal Modes . . . 123
6.4.2 Bandwidth Conversion . . . 124
6.5 Optimising the Raman Light-Matter Interaction . . . 127
6.6 Summary . . . 129

7 Conclusions and Outlook 131
7.1 Summary . . . 131
7.2 Future Work . . . 132
7.2.1 Further Noise Suppression . . . 132
7.2.2 Temporal Multiplexing . . . 133
7.2.3 Light-Matter HOM . . . 133
7.2.4 Temporal Mode Verification . . . 133
7.2.5 High-Efficiency Single-Mode Cavity Memory . . . 134
7.3 Conclusions . . . 134

Bibliography 135

A Green's Function analysis 150

List of Figures

1 Quantum Memories Group Picture 18

1.1 Schematic of a Basic Quantum Repeater Protocol 28
1.2 Temporal Multiplexing with Quantum Memories 29
1.3 Energy Level Schematics for Optically-Controlled Memory Protocols 33

2.1 Raman Memory Protocol 37
2.2 Schematic of the Raman Memory Protocol and Four-Wave Mixing Noise 38
2.3 Suppression of Four-Wave Mixing Via Selection Rules 52

3.1 Schematic of the Raman memory 55
3.2 Experimental Set-up of the Raman Memory 56
3.3 Spin Polarisation as a Function of Time 58
3.4 Frequency Calibration using Saturated Absorption Spectroscopy 59
3.5 Set-up of the Frequency Filtering Etalons 62
3.6 Measured Spin-Polarisation of Caesium in Different Buffer Gases 64
3.7 Typical Raman Memory Trace in High Gain Regime 65
3.8 Efficieny Vs Rabi Frequency at Different Temperatures 66
3.9 Efficieny Vs Rabi Frequency at Different Detunings from Resonance 67
3.10 High Gain Raman Memory 68

4.1 Ring Cavity Containing Caesium Atoms 71
4.2 Cavity Triple Resonance Condition 76
4.3 Experimental Set-up of the Cavity Memory 78
4.4 Cavity Resonance Conditions 79
4.5 Second-Generation Cavity Resonance Conditions 81
4.6 Experimental Set-up for Second-Generation Cavity 82
4.7 Transmission of a Pulse On-Resonance with the Cavity 83
4.8 Scaling of Noise in a Cavity-Enhanced Raman Memory 84
4.9 PDC Spectrum from PPKTP Waveguide Source Vs Temperature 87
4.10 Feed Forward Schematic 88
4.11 Single Photon Storage in a Cavity-Enhanced Raman Memory 89

5.1 Raman Memory with Built-In Noise Suppression Schematic 93
5.2 Numerical Simulations of the Built-In Noise Suppression in the Raman Memory . 96
5.3 Four-Wave Mixing Phase Matching as a Function of Detuning 98
5.4 Simulated Memory Efficiency Vs Detuning for Different Control Powers 98
5.5 Simulated Memory Efficiency Vs Detuning for High Optical Depth 99
5.6 Experiemental Set-up for Absorptive Raman Memory 100

5.7 Memory Efficiency as a Function of Control Energy for BNS-Raman Memory . . 101
5.8 Noise Vs Control Power and Lifetime for BNS-Raman Memory 102
5.9 Noise in Multiple Read-Out Bins 102
5.10 Coincidence Traces to Measure Autocorrelation Functions 104
5.11 $g^{(2)}$ as a Function of N_{in} for the BNS Raman Memory 104
5.12 Noise as a Function of Optical Pumping Efficiency 106
5.13 Noise as a Function of Pulse Duration 107
5.14 Predicted $g^{(2)}$ for a Single Photon Input 108

6.1 Temporal Modes in Time-Frequency Space 111
6.2 Pulse Carver 114
6.3 TA Characterisation 116
6.4 Pulse Distortion due to High TA Current 117
6.5 Schematic of Set-up for Temporal Modes Memory 118
6.6 Mode Selectivity of the Raman Memory 119
6.7 Tomography of Raman Memory 121
6.8 Mode Conversion between Hermite-Gaussian Temporal Modes 124
6.9 Hermite-Gaussian Mode Conversion for Different Control Energies 125
6.10 Bandwidth Conversion using the Raman Memory 126
6.11 Bandwidth Conversion Efficiency 126
6.12 Conversion to Falling Expoinentials 127
6.13 Optimising the Memory Efficiency via Tailored Temporal Modes 128
6.14 Optimising the Output Temporal Mode of the Memory 129

List of Tables

5.1 Fit parameters for $g^{(2)}$ of the retrieved state as a function of photon number . . 106

6.1 Extracted mean and limiting cases of the single-modeness for a projection of $\boldsymbol{P}$ onto controls in a 3-dimensional subspace. 122

Abbreviations, Symbols and Constants

List of Abbreviations

AFC Atomic Frequency Comb

AOM Acousto-Optic Modulator

APD Avalanche Photodiode

ASE Amplified Spontaneous Emission

AWG Arbitrary Waveform Generator

BNS Built-in Noise Suppression

BS Beam-Splitter

BSM Bell-State Measurement

CRIB Controlled Reversible Inhomogeneous Broadening

CW Continuous Wave

DFB Distributed Feedback Diode Laser

ECDL External Cavity Diode Laser

EIT Electromagnetically-Induced Transparency

EOM Electro-Optic Modulator

FLAME Fast Ladder Memory

FMZI Fibre Mach-Zehnder Interferometer

FPE Fabry-Pérot Etalon

FSR Free Spectral Range

FWM Four-wave mixing

GEM Gradient Echo Memory

GTI Giles-Tournois Interferometer

HeNe Helium-Neon Laser

HG Hermite-Gauss (polynomials)

LOQC Linear-Optical Quantum Computing

OAM Orbital Angular Momentum

ORCA Off-Resonant Cascaded Absorption

OSA Optical Spectrum Analyser

PBS Polarising Beam-Splitter

POVM Positive Operator Value Measure

PPKTP Periodically-poled potassium-titanyl-phosphate

QD Quantum Dot

QIP Quantum Information Processing

QKD Quantum Key Distribution

QM Quantum Memory

QPG Quantum Pulse-Gate

REID Rare-Earth Ion-Doped crystal

RF Radio Frequency

SFPI Scanning Fabry-Pérot Interferometer

SHG Second Harmonic Generation

SMF Single-Mode Fibre

SNR Signal to Noise Ratio

SPDC Spontaneous Parametric Down-Conversion

SPIDER Spectral Phase Interferometry for Direct Electric-field Reconstruction

SQL Standard Quantum Limit

SRS Spontaneous Raman Scattering

SVD Singular Value Decomposition

TA Tapered Amplifier

TEAPARTI Temporal Electric-field Amplitude & Phase: Algorithmic Reconstruction via Time-domain Interferometry

Ti:Sapph Titanium Sapphire oscillator

TM Temporal Mode

List of Symbols and Constants

Fundamental Constants

c ..... Speed of light

ε_0 ..... Permittivity of free space

μ_0 ..... Magnetic permeability of free space

$\hbar$ ..... Reduced Planck constant

k_B ..... Boltzmann constant

Mathematical

$*$ ..... Complex conjugate

$\dagger$ ..... Hermitian conjugate

$\boldsymbol{\nabla}$ ..... Nabla gradient operator

$\otimes$ ..... Tensor product

i ..... Imaginary unit

$\mathbb{1}$ ..... Identity operator

$\mathfrak{Im}\{z\}$ ..... Imaginary component of z: $\mathfrak{Im}\{z\} = \frac{1}{2\mathrm{i}}(z - z^*)$

$\mathfrak{Re}\{z\}$ ..... Real component of z: $\mathfrak{Re}\{z\} = \frac{1}{2}(z + z^*)$

$\mathrm{Tr}[\rho]$ ..... Trace of ρ

Definitions for Quantum Memories

Λ ..... "Lambda" configuration of a three level system

Δ ..... Single photon detuning

Δ_hf ..... Hyperfine ground state splitting

$\delta_{c,s}$ ..... Bandwidth of control/signal field

ω ..... Angular frequency

L ..... Length of atomic medium

d ..... Resonant optical depth

η_mem ..... Memory efficiency

μ_1 ..... Metric for noise level

τ_mem ..... Memory lifetime

B ..... Time-bandwidth product

$g^{(2)}$ ..... Second order autocorrelation function

κ ..... Single-modeness

List of People Directly Acknowledged

B.B. Dr Benjamin Brecht

D.J.S. Dr Dylan Saunders

J.H.D.M. Dr Joseph Munns

J.K. James Klatzow

J.N. Dr Joshua Nunn

J.S. Dr Jan Sperling

O.L.A. Oscar Lazo-Arjona

P.M.L. Dr Patrick Ledingham

T.M.H. Thomas Hird



Figure 1: The quantum memories team

Chapter 1:

Introduction

1.1 Quantum-Enhanced Technology

Quantum technologies harness the power of quantum mechanics and exploit features such as superposition, entanglement and interference to increase performance over classical counterparts. The applications of quantum technologies include precise sensing and measurement [1], secure long-distance communication [2], simulations of currently intractable systems [3] and exponentially faster computations [4]. This exciting and diverse field has gained huge momentum over the past few decades, and I will give a brief outline of just some of the extensive applications and ongoing research in the following sections.

1.1.1 Quantum Computation and Simulation

The ultimate goal for many in this field is to produce a "universal quantum computer" which consists of quantum bits, or "qubits", on which one can perform logical operations and gates. In analogy to a classical computer in which information is encoded in bits which can either be a zero or a one, a qubit can exist in two states representing a zero or a one or any superposition of those qubit states. An array of n qubits can therefore be in an arbitrary superposition of 2^n states. A quantum algorithm is performed by preparing the qubits in a particular initial state, performing a series of logical operations, and then measuring the final state which collapses the system into one of the 2^n eigenstates [5]. It has been shown that any quantum algorithm can be represented as a quantum circuit, or a sequence of quantum gates, which consists of only single- and two-qubit operations, and that this is sufficient for universal quantum computation [6]. Since qubits can exist in a superposition of states, an array of qubits can explore a much larger set of possible solutions and hence quantum algorithms can solve certain problems exponentially faster than the best known classical algorithms [7–9].

In addition to this circuit model of quantum computation described above, there are alternative routes to universal quantum computation including measurement-based protocols, or

"one-way quantum computing". Here a large entangled cluster state is generated and universal computation can be performed using single qubit operations and measurement [10]. Classical feed-forward is used to determine which measurements should be performed given the results of earlier ones, which encodes the quantum logic gate on the state. The advantage of one-way quantum computing is that two-qubit operations are not required, although the generation of large cluster states is experimentally challenging.

A secondary approach to quantum computation is using a quantum system to simulate a different physical quantum system, an idea first proposed by Feynman in 1982 [3]. A quantum simulator, or analogue quantum processor, cannot necessarily perform universal computations but can be used to calculate information about complex systems [11]. An adiabatic quantum simulator can be used to find the ground state of a Hamiltonian which describes a system of interest, for example the stable configuration of a complex molecule [12, 13]. The system is prepared in the ground state of a known, simple Hamiltonian, and then the Hamiltonian is adiabatically evolved by slow adjustment of experimental parameters to the desired Hamiltonian. Here "slow" is compared to the energy gap between eigenvalues of the Hamiltonian, such that the system remains in the ground state. The final state of the system is the ground state of the Hamiltonian of interest, and hence the energy of this ground state can be calculated. It has been shown that NP-hard problems [14] can be efficiently reduced to the problem of finding the ground state of a Hamiltonian and therefore can be embedded on adiabatic quantum computers, provided there is full connectivity between the qubits [15]. This "quantum annealing" method can therefore be used to solve optimisation problems, such as finding the optimal radiation dose distribution to kill cancer cells whilst minimising damage to healthy tissue [16], training of deep neural networks for machine learning [17], and simulation of complex materials such as topological states of matter [18].

The potential of quantum computers to solve problems that are intractable on today's classical computers has led to extensive experimental research, and I will discuss some of the platforms and recent advances below in Section 1.2.

1.1.2 Quantum Communication

In addition to quantum computation, a large part of quantum technologies focuses on quantum communication which can be split into two closely-related objectives: being able to faithfully transfer a quantum state over large distances, and having unconditionally secure communication channels.

Quantum mechanics is subject to the no-cloning theorem, which states that it is impossible to deterministically copy an arbitrary unknown quantum state [19]. It is hence not possible to amplify a quantum signal to make it robust against a lossy channel without adding noise to the state, and therefore there is a limit on the distance over which one can faithfully transmit a quantum state. One way to transmit quantum information from one location to another is to use quantum teleportation [20] which relies on a shared entangled state and a classical communication channel between parties. The problem of transferring quantum states can be reduced to the ability to share an entangled state over large distances, beyond the limit of 10s of km that can be reached via direct, lossy transmission, and this is a point we will come back

to in Section 1.3.1(a).

In 1994 Shor published a quantum algorithm that could perform prime factorisation of large integers exponentially quicker than the best known classical algorithm [7]. This has the potential to break current secure communication protocols, which rely on the fact that it is intractable to factorise large numbers on classical computers. Shor demonstrated that a large quantum computer would be able to defeat cryptography schemes, which has led to extensive research into quantum cryptography.

Quantum cryptography provides the means to secure a communication channel against an eavesdropper, even if they have access to a quantum computer. The most common application is the distribution of a secret key between two parties, or quantum key distribution (QKD). In QKD a quantum key is encoded in a random string of qubits and by discussing publicly what basis was used to prepare and measure the qubits, the users can share a secret key whilst establishing whether an eavesdropper was present [2]. QKD relies on the no-cloning theorem of quantum mechanics since an eavesdropper cannot copy the qubits in a random cryptographic key and would always perturb their state by measuring them, alerting the users to their presence. Once a quantum key has been successfully shared, a secure classical communication channel can be established. Common QKD protocols rely on shared entanglement between the users [21], and therefore we again see that it is critical to be able to distribute entanglement over large distances.

1.1.3 Quantum Metrology

A final branch of quantum technologies that we will briefly discuss is quantum metrology, which is the field of high-resolution measurements and parameter estimations [22, 23]. Quantum metrology uses quantum effects, in particular entanglement and squeezing, to enhance the precision of measurements beyond what is possible using classical approaches. The central limit theorem of classical estimation theory states that the statistical error in a measurement scales as $n^{-1/2}$, where n is the number of times the measurement is repeated, and this scaling is known as the "standard quantum limit" (SQL). Using quantum effects such as entanglement makes it possible to beat the SQL and approach the fundamental limit set by Heisenberg uncertainty relations where the error scales as n^{-1} [1]. If there are interactions between the particles used to perform the measurements then even better scaling, referred to as "super-Heisenberg", can be achieved [24, 25].

1.2 Architectures for Quantum Technologies

The exciting and varied applications of quantum information processing (QIP) has led to extensive research into a myriad of different physical platforms. The core requirements for QIP are summarised by the DiVincenzo criteria [26]:

1. A scalable physical system with well characterized qubits
2. The ability to initialise the state of the qubits
3. Long relevant decoherence times

4. A "universal" set of quantum gates
5. A qubit-specific measurement capability
6. The ability to interconvert stationary and flying qubits
7. The ability to faithfully transmit flying qubits

where the final two are necessary for quantum communication, and not for quantum computation.

There are many different potential architectures for quantum information processing each with their own distinct advantages and drawbacks, and each satisfying a different subset of these criteria. The most promising of these thus far for quantum computation are based on single trapped atoms [27] or ions [28], superconducting circuits [29], silicon-based platforms [30], and solid-state defects [31], whilst the natural platform for quantum communication is undoubtedly photonics [32]. We discuss each of these platforms and their pros and cons below.

1.2.1 Trapped Ions

Atomic or molecular ions can be isolated and spatially trapped using electromagnetic fields. The internal electronic state of an ion can be used as a qubit, and preparation and measurement of qubits can be performed optically, whilst single-qubit gates can be realised by optical or microwave addressing of individual ions. Two-qubit gates are performed by coupling the internal qubit states to collective external motional states of the array of ions. Single- and two-qubit gates have been demonstrated with fidelities of 99.9934(3)% and 99.9(1)% respectively [33]. The speed of two-qubit gates is limited by the characteristic frequency of the motional states of the ions and, until recently, has been limited to hundreds of microseconds [33, 34] which significantly limits the computational speed. However, recent work overcomes this limitation by dynamic shaping of the laser pulses driving the interaction, and two-qubit entangling gate times of 1.6 μs with a fidelity of 99.8% have been demonstrated [35]. The coherence time of the qubits is on the order of one second, and is typically limited by stray magnetic fields. The largest number of ions to be confined and fully controlled in a single trap is 20 [36], and it is challenging to increase the number of ions whilst maintaining the ability to optically address individual qubits. There are proposals to reach scalability of this platform by using an array of small computational nodes which are connected by photonic links. However, despite the small number of qubits, quantum simulations to calculate the ground state energy of molecules have already been performed using four trapped ions [37].

1.2.2 Single Atoms

An alternative platform for QIP uses the internal states of a single atom to encode information, which is again an attractive platform due to the long coherence times and ability to cool and trap atoms in an optical lattice [38]. By exciting the trapped atoms to Rydberg states, neighbouring atoms interact via strong dipole-dipole interactions, and fast single-qubit [39] and two-qubit gates [40, 41] can be performed. Fidelities of 99.62(16)% for single-qubit gates [42] and 81(2)% for two-qubit gates [43] have been demonstrated, and by careful control of the trapping magnetic and optical fields, coherence times of a minute have been reached [44]. A quantum simulator with

an array of 51 trapped Rydberg atoms has been realised to perform simulations of symmetry-breaking phase transitions [45].

1.2.3 Superconducting Circuits

Superconducting quantum computers use an array of Josephson junctions which act as artificial atoms with anharmonic energy levels, and a comprehensive review is given in [46]. Superconducting electronic circuits can be written on the micrometer scale using established integrated circuit technology, and need to be cooled below 100 mK. There are several main types of superconducting qubits including charge qubits, where the qubit states correspond to different numbers of Cooper pairs within the junction, and flux qubits which correspond to different integer numbers of magnetic flux quanta. The energy gap between the qubit states is on the GHz scale and therefore single qubit gates are performed via microwave pulses. Two-qubit gates are performed by coupling distinct qubits to an intermediate electrical component such as a capacitor. Single- and two-qubit gate fidelities of 99.92% and 99.4% respectively have been achieved [47], and typical gate times are 130 ns for single-qubit and 250 – 450 ns for two-qubit gates [34]. The coherence times of superconducting qubits are on the order of 120 μs [48]. Since superconducting qubits are fabricated on chip, this platform is readily scalable and large 2D arrays of qubits have been fabricated. The IBM Quantum Experience enables users to run quantum algorithms on a 16 qubit processor [48], and testing of larger circuits has been announced by IBM [49], Intel [50] and Google [51]. D-wave Systems have commercial systems with over a thousand superconducting qubits [52], but these devices can only perform quantum annealing simulations rather than universal quantum computing.

1.2.4 Solid State Protocols

A scheme for scalable quantum computing using phosphorus donor atoms embedded in a silicon lattice was proposed by Bruce Kane [30]. Information is encoded in the nuclear spins of the phosphorus atoms, which are well isolated from the environment leading to very long decoherence times of over one second for a small-scale device at 5.6 K [53], and over half an hour in isotopically pure bulk silicon [54]. Single-qubit operations are performed by applying external electric fields to metallic gates deposited on the semiconductor surface above each donor. Two-qubit gates are performed by transferring the spin information from the nucleus to the donor electron, and then using a gate voltage to draw donor electrons from adjacent spins to an intermediary region where they interact with a much higher coupling strength. Two-dimensional arrays of qubits have been produced, with an electrode pattern above and below the qubit layer that allows for control and read-out of individual qubits with a fidelity of 97.9% [55]. Whilst high-fidelity qubit gates have not yet been demonstrated, the major advantage of this platform is the well-established silicon engineering industry from classical computers, which leads to a feasible scalability of this technology that other technologies such as trapped ions perhaps lack.

A secondary platform using silicon technology is based on quantum dots (QDs), which are nano-scale semiconductors that act as artificial atoms. QDs were first proposed as a platform for quantum computation in 1998 [56], and initial research focused on III-V semiconductor compounds such as GaAs [57, 58]. Initialisation and measurement of the qubits can be performed

using electrical pulses, and two-qubit gates are achieved by electrical gating of the tunnelling barrier between neighbouring QDs. However, the nuclear spin background of the semiconductor leads to strong dephasing and a limited coherence time of the qubits [58]. A significant improvement in the coherence time can be found by using quantum dots in silicon [59, 60], which can be isotopically purified to give dephasing times of 61 μs [61]. Two-qubit gates have been performed [61], but do not yet reach the fidelity level of trapped ions or superconducting qubits.

Alternative approaches using spins in solid state systems include using defects such as nitrogen-vacancies in diamond [62], divacancy defects in silicon carbide [63], and rare earth ions in solids which offer six hour coherence times [64]. A comprehensive review of solid state architectures for quantum technologies is given in [31].

1.2.5 Hybrid Platforms

The results presented above demonstrate the impressive progress that is being made across a variety of platforms, but scalability still remains a significant challenge. It is possible that the route towards scalable quantum technologies will be to harness the unique advantages of different platforms in a hybrid system. We therefore turn our attention to quantum networks or a "quantum internet" scheme [65], where small, few-qubit computational nodes are connected via photonic links in a modular approach. The computational nodes will allow high fidelity, deterministic gates to be performed, and information is transferred to flying qubits, or photons, using a light-matter interface. Photonic networks are a key element of scalable quantum technology, and the focus of this thesis. In the next section we discuss photonic QIP in more detail.

1.2.6 Photonic Quantum Information Processing

Optical QIP offers a very promising platform for quantum communication and for connecting nodes of a quantum network, as we discussed above. In addition, there are proposals for all-optical QIP schemes, which some argue will be the easiest platform to scale into large, fault-tolerant quantum computers in the near future [66]. There are several key advantages to photonic QIP. Firstly, photons are essentially free from thermal noise at optical frequencies at room temperature, and therefore it is easy to produce very low noise qubits. Secondly, photonic integrated circuits are a well-established technology in classical computation and communication, and a lot of the relevant engineering and infrastructure necessary for photonic circuits and networks is already in place. Finally, light has a large information capacity due to its large bandwidth [67], and photons have many degrees of freedom in which to encode information, including polarisation [68], time-bins [69], frequency-bins [70], temporal modes [71], and spatial modes [72], and it is straightforward to perform single-qubit gates via polarisation rotations, pulse delays, or spatial light modulators.

However, the fact that photons interact very weakly means it is very challenging to induce strong photon-photon interactions to perform two-qubit gates which are critical for the circuit model of universal quantum computation. There are two main approaches to solve this. Firstly, it has been shown that linear optical operations, using beam splitters, mirrors and phase shifters, in addition to single photon sources, detectors and ancilla qubits, are sufficient

for a full photonic quantum information protocol, known as linear-optical quantum computing (LOQC) [73]. However, these linear optical gates are inherently probabilistic which prevents the technology from being effectively scalable. Secondly, research into deterministic single-photon non-linearities using nonlinear optics [74], Rydberg atoms [75] and quantum dots [76], has seen great progress, but it remains a significant challenge to induce a large enough phase shift at the single photon level.

In addition to high fidelity single- and two-qubit gates, a vital component for photonic QIP is a bright, deterministic source of pure, indistinguishable single photons. The metric that is used to quantify whether a photonic state is a true single photon is the Glauber second-order autocorrelation at zero time delay, $g^{(2)}$, which is a measurement of the probability of two simultaneous photons, and a true single photon state will have $g^{(2)} = 0$ [77]. It is also crucial that the photons are in a pure, rather than mixed, state and in particular we need the photons to be in a single well-defined temporal and spectral mode. Furthermore, every photon must be identical and we use the visibility of interference measurements, V, to characterise the indistinguishability.

Significant advances have been made towards high quality single photon sources across many different platforms. The most commonly used single photon sources use nonlinear optical effects to generate pairs of photons, for example spontaneous parametric down conversion (SPDC) sources where a high energy pump photon is converted into a pair of lower energy photons, and the detection of one down-converted photon heralds the presence of the other. More details about SPDC are given in Chapter 4. Photon sources based on parametric down conversion were first demonstrated in 2001 [78], and have since seen significant advances from controlling their spectral-temporal modes [79], generating photons at telecommunication wavelengths [80], and operating sources in a nonlinear waveguide [81] which allows high coupling efficiencies into optical fibres [82]. Today SPDC sources can produce pure single photons with very high purity, $g^{(2)} \sim 0.002$ [83], high indistinguishability $V \gtrsim 99\%$ [84], and heralding efficiencies of over 80% [85, 86]. An alternative approach to heralded single sources uses the nonlinear optical effect of spontaneous four-wave mixing (SFWM) in optical fibres [87] and in waveguides [88]. This approach allows the generation of multiple, near-identical sources on the same silica chip, which increases the scalability of this technology.

The bandwidth of the photons generated via SPDC or SFWM is determined by the phase-matching condition which depends on the dispersion of the medium, the spectral bandwidth of the pump and the inverse length of the interaction medium, and typically the bandwith of the downconverted photons is on the order of several nanometres. However, in order to have an efficient light-matter interface to transfer information from matter-based stationary qubits to photonic flying qubits, we require photonic bandwidths on the MHz–GHz scale, several orders of magnitude less than down-converted photons. The generated photons can be filtered down to atomic-compatible bandwidths [89], but this dramatically reduces the photon generation rate to 10s of kHz. Fibre-based sources offer long interaction lengths to generate narrower photon pairs [90, 91], but the loss in the fibre becomes an issue and limits the achievable heralding efficiency. Significant advances have been made to operate SPDC sources inside a cavity to produce sub-MHz bandwidth photons [92], and in a resonant waveguide to produce photons

with tunable bandwidths on the MHz–GHz scale [93].

However, in both SPDC and SFWM sources the pair production probability must be kept low to prevent multi-photon events, and therefore there is an intrinsic compromise between brightness and $g^{(2)}$. The largest entangled state of photons generated thus far is a 12 photon GHZ state [94], which was generated using six heralded single photon sources. However, the generation rate here was approximately one per hour, which highlights the scaling problem of probabilistic photon sources.

Recently, much progress has been made on a promising alternative approach: single photon sources from single emitters. Early demonstrations focused on atoms [95], ions [96] and molecules [97], but the most promising sources today are based on quantum dots [98]. A quantum dot (QD) is a nano-scale semiconductor embedded in a higher band-gap semiconductor which acts as an artificial atom, and QD sources do not suffer from multi-photon events like SPDC sources as an isolated QD is a single emitter. The most promising QD sources are based on III-V semiconductors such as InAs and GaAs [99, 100] which emit at infra-red frequencies. QDs can be inserted in micropillar [101, 102], microdisk [103], and photonic-crystal cavities [104] to increase the spontaneous emission rate due to the Purcell effect and reach the strong-coupling regime. Micropillar cavities can be deterministically grown around a single QD to enhance emission into a single spatial mode and allow high collection efficiencies [105, 106], and very bright sources with an emission rate of 11 MHz have been measured [107]. However, fluctuations of charges in the nearby environment of the QD lead to frequency jitter of the emission and therefore it is challenging to have a high indistinguishability source which emits into a pure spectral mode. Alternatively QDs can be grown via aluminium droplet etching and placed under solid immersion lenses to enable high interference visibilities of 95% without the need for Purcell enhancement [108], although the collection efficiency of these devices is not yet as high.

State of the art sources of pure, indistinguishable photons and photonic gates both rely on probabilistic processes, and hence there is an inherent scaling problem in photonic QIP. As the size of a photonic network increases, the overall success probability becomes vanishingly small. Therefore, a vital requirement for photonic quantum technology is the ability to store a quantum state of light, which would allow multiplexing of probabilistic operations, in a "repeat until success" protocol. A quantum memory is a device that offers this capability, and is described in detail below.

1.3 Quantum Memories

An optical quantum memory is a device that can coherently store and retrieve, on-demand, an optical quantum state without destroying or disrupting its quantum state. Quantum memories map the travelling optical signal into a stationary medium, and the process needs to be coherent such that the non-classicality is preserved for retrieval at a later time. By "on-demand" we require that the time at which the state is retrieved from the memory can be dynamically chosen, after the state has been stored.

Optical quantum memories can be used to locally synchronise probabilistic gates and opera-

tions at a computational node, synchronise probabilistic single photon sources using a temporal multiplexing strategy, and to distribute quantum information over large scales in a quantum network. We discuss these applications in Section 1.3.1, outline the figures of merit that are required from a quantum memory in Section 1.3.2, and finally discuss the different experimental implementations of optical quantum memories in Section 1.3.3

1.3.1 Applications of Quantum Memories

1.3.1(a) Quantum Repeaters

We discussed in Section 1.1.2 that sharing entanglement between users is vital for quantum communication, both for teleporting quantum states from one location to another and for transmitting a secret key between users to enable secure long-distance communication. However, there is a maximum distance over which one can send a quantum state due to loss in optical fibre. In classical communication light can travel hundreds of kilometres in optical fibres, and optical amplifiers are used to amplify the light signal along the length of the connection. These amplifying stations are known as "repeaters" and extend the range of optical communication to global length scales. However, qubits are subject to the no-cloning theorem and thus cannot be deterministically amplified without adding noise. Fibre losses are typically around 0.2 dB/km [109], and therefore after 100 km 99% of a signal would be lost.

A "quantum repeater" is a device that can be used to extend the range of quantum communication by enabling the distribution of entanglement over large distances [110]. The basic principle of a quantum repeater is to subdivide the transmission link into multiple short segments over which entanglement can be faithfully distributed, and then the entanglement range is increased using a process known as "entanglement swapping" between neighbouring segments [111]. Pairwise entanglement is established between nodes 1 and 2, and 3 and 4, as shown in Figure 1.1. A Bell-state measurement is then performed between nodes 2 and 3, which yields entanglement between the further-separated nodes 1 and 4. These entanglement swapping operations are performed in a nested fashion until the entanglement is distributed across the entire communication link. This enables long-distance entanglement with a polynomial scaling in time.

However, typically the Bell-state measurements (BSM) are performed using linear optical operations and are probabilistic, and hence the probability of successfully swapping entanglement at every node of the communication link scales poorly with the number of nodes. This means that quantum memories are a vital component of quantum repeaters to synchronise the creation of adjacent links and successful entanglement swap operations. The quantum memory can store the entangled states whilst waiting for other BSMs to herald successful entanglement between neighbouring nodes.

The pioneering approach for the experimental realisation of a quantum repeater is referred to as the DLCZ protocol, after the authors of the original paper (Duan, Lukin, Cirac and Zoller) [113]. Here, an entangled state of a photon and an atomic excitation is generated via Raman scattering. The photons emitted from two DLCZ memories are interfered on a beam splitter, and the detection of one photon generates entanglement between the two atomic

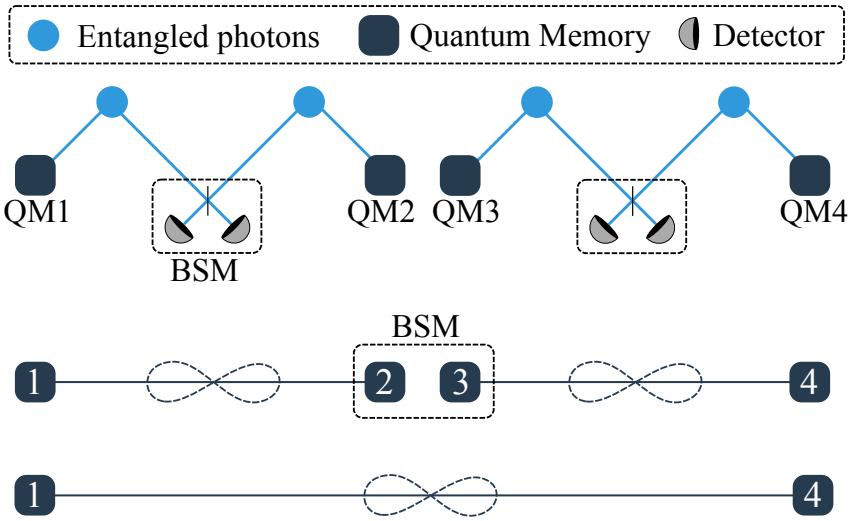


Figure 1.1: Schematic of a basic quantum repeater protocol using quantum memories. One photon from an entangled pair is stored in a quantum memory, and the other is interfered on a beam splitter with a photon from a neighbouring source. The detection of one photon at the output of the beam splitter yields entanglement between the two quantum memories. This is equivalent to a Bell state measurement (BSM) between the two sources. If entanglement is established between memories 1 and 2, and 3 and 4 respectively, then the entanglement can be extended by performing a BSM between memories 2 and 3. This results in entanglement between memories 1 and 4. This scheme is a basic version of a quantum repeater protocol and significantly more sophisticated versions now exist, for a review see [112].

ensembles, since it is not known from which ensemble the photon originated. Entanglement can then be extended by performing BSMs between further pairs of entangled ensembles.

The necessity of quantum repeaters for long distance quantum communication has been the driving force behind the field of quantum memories. However, there are many other applications for quantum memories including single photon detection [114, 115], quantum metrology [116, 117], and, crucially for LOQC, local synchronisation [118].

1.3.1(b) Local Synchronisation

We discussed above in Section 1.2.6 that LOQC relies on probabilistic gates, and also that commonly-used single photon sources rely on probabilistic processes such as spontaneous parametric down conversion. In order to make LOQC schemes scalable, it is vital to have quantum memories to synchronise these probabilistic events and increase the overall success probability of photonic networks. The successful outcome of a probabilistic operation can be temporarily stored until other nearby processes have also succeeded, in a "repeat until success" temporal multiplexing strategy.

We use the example of SPDC photon sources to demonstrate this idea of temporal multiplexing using quantum memories. If we have N probabilistic single-photon sources and N quantum memories, then whenever a photon is generated from a source it can be stored in the corresponding memory until $N-1$ sources have produced photons. When the final source generates a photon, the photons can be released from the memories synchronously in a $N-$photon state, as shown in Figure 1.2. This dramatically reduces the waiting time for multi-photon

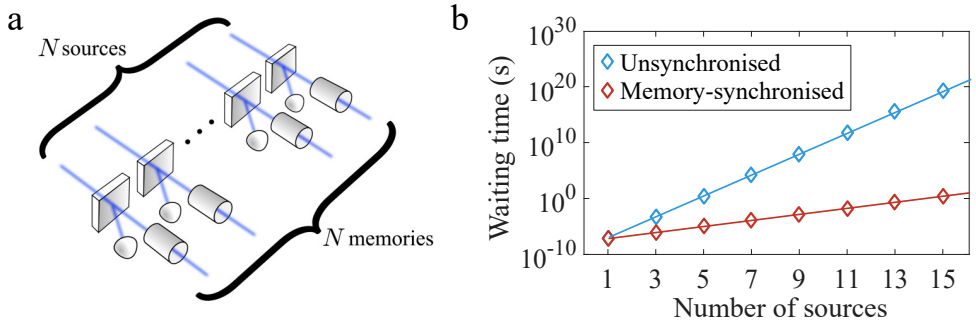


Figure 1.2: (a) A schematic showing temporal multiplexing of N probabilistic single photon sources using N quantum memories. When the generation of a photon from one of the sources is heralded by detection of the sister photon, the photon is stored in the memory. When $N-1$ memories are loaded, and the final source generates a photon, all the photons are retrieved from the memories simultaneously. Figure from [118] (b) The predicted waiting time for a N-photon state from N unsynchronised sources increases exponentially. If quantum memories are used to temporally multiplex the sources, the average waiting time significantly decreases. The memories used in the simulation here have a time bandwidth product of $B = 1000$, an efficiency of $\eta = 56\%$ and a repetition time of 1 GHz.

states and overcomes the scaling problem of increasing the size of a photonic network [118].

The quantum memories used to synchronise local operations in this way do not necessarily need to have long storage times. Instead, we require a large synchronising capacity, or a large storage bandwidth of the memory with respect to its lifetime to have many storage attempts within the lifetime of the memory. This is in contrast to memories for quantum repeaters which need to have long storage times with respect to the distance over which we want to distribute entanglement. The requirements of a quantum memory depend on its application, and it is likely that different platforms will be used for these varying applications. We discuss the desired properties of quantum memories more in the following section.

1.3.2 Figures of Merit

Quantum memories have a broad range of applications from long-distance entanglement distribution to local synchronisation, and the necessary performance criteria for quantum memories varies greatly between applications [119]. A summary of the key figures of merit for quantum memories is given below.

- **Efficiency**
  The memory efficiency, η_{mem}, is defined as the ratio of the energy of the input pulse and the retrieved pulse. High memory efficiency is a desired property for all quantum memory applications. For practical applications we also require a high end-to-end efficiency, which includes the internal memory efficiency as well as losses from the memory protocol itself including optical components and filtering.

- **Fidelity**
  We require low noise memories such that the retrieved state has a high fidelity with the input state and photon number statistics are preserved. The fidelity between the input state, ρ_{in}, and the output state, ρ_{out}, is defined as

$$F(\rho_{\text{in}}, \rho_{\text{out}})) = \text{Tr}\left[\sqrt{\sqrt{\rho_{\text{in}}}\rho_{\text{out}}\sqrt{\rho_{\text{in}}}}\right]^2 \tag{1.1}$$

and quantifies the similarity between the two states. We ideally want a unit fidelity, $F = 1$, for any input state, ρ_{in}. The threshold fidelity to beat a classical measure-and-prepare strategy is $F > 2/3$ for a single photon qubit [120], and for practical applications we require a high enough fidelity to reach the fault-tolerant quantum error correction threshold [121].

A useful metric we use to characterise the noise level in a quantum memory is $\mu_1 = \langle n_{\text{noise}}\rangle/\eta_{\text{mem}}$, where $\langle n_{\text{noise}}\rangle$ is the average number of noise photons per memory trial, and μ_1 quantifies the noise-to-signal ratio on retrieval [122]. The μ_1 parameter allows us to compare noise between different quantum memories, but the true metric of whether the memory is noise-free is to measure the fidelity [119]. For the case of storing single photons one can measure the Glauber autocorrelation of the output and see if the photon number statistics are unchanged [123], which gives an upper bound on the fidelity for the storage of qubits.

- **Lifetime**

  We want an "on-demand" quantum memory, where we can choose the time at which the state will be retrieved from the memory after it has been stored. The stored atomic excitation will decohere due to various dephasing mechanisms, and this limits the timescale over which the state can be stored. In this thesis we define the memory lifetime, τ_{mem}, as the time at which the memory efficiency has decreased to $1/e$ of its initial value, although other definitions include the storage time at which the fidelity drops below the classical benchmark [124]. For quantum repeater protocols the storage time needs to be at least as long as the average time for entanglement generation which is typically on the order of 1 second [112], but can be decreased if the entanglement generation is parallelised [125]. There are quantum repeater protocols that do not directly use quantum memories [126], but these require large entangled states which are built using local synchronisation. The lifetime requirements are significantly less stringent for quantum memories used for local synchronisation.

- **Bandwidth**

  The bandwidth of a quantum memory, δ, is defined as the maximum bandwidth of an optical signal that can be stored. For local photonic networks we want to perform fast operations and therefore require high bandwidth photonic states, and memories for local synchronisation applications need a high storage bandwidth. Bandwidth considerations are also important when considering interfacing single photon sources with atomic memories, as we discussed in Section 1.2.6.

- **Time-bandwidth product**

  The time-bandwidth product is defined as $B = \tau_{\text{mem}}\,\delta$. B quantifies the maximum number of storage attempts within the lifetime of the memory and is therefore a measurement of

the synchronising capacity. For temporal multiplexing of probabilistic sources and gates, $B \gtrsim 1000$ offers a dramatic advantage [118].

- **Modal Capacity**
  We define the modal capacity, N_{modes}, as the number of modes that the memory can efficiently store, whether that be spatially [127, 128], spectrally [129, 130] or temporally [131, 132]. If we have a memory that can store N modes then the quantum repeater scheme can be multiplexed to achieve an N-fold increase in overall success probability [133], and thus a large multimode capacity of memories is an appealing feature for quantum repeaters. On the other hand, if a memory can only store a single mode then it can be used to operate on temporal mode qubits and enhance the purity of single photon sources [134]. The desired modal capacity depends of the application of the quantum memory.

- **Scalability**
  All the applications of quantum memories that we have discussed require a large number of devices, and therefore we need the memory itself to be a feasibly scalable platform. Requirements such as room temperature operation, convenient operational wavelengths and small technical overheads are therefore appealing attributes.

1.3.3 Experimental Implementations of Quantum Memories

In the last few years there has been significant progress towards experimental realisations of optical quantum memories, and extensive reviews are given in [119, 124, 135]. I will give a brief outline of the different approaches in the following sections, although this overview is by no means exhaustive.

Memory protocols can be split into two distinctive categories: emissive quantum memories, where the photon is generated in the memory itself, and absorptive memories, which can store a quantum state from an external source. Emissive memories, such as ones in the DLCZ protocol, can be thought of as a photon-source and memory combined, and the generated photons are automatically compatible with the atomic wavelength and bandwidth. However, to enable synchronisation of probabilistic logical operations and single photon sources we require an absorptive quantum memory which offers storage of other photonic states.

The simplest realisation of an absorptive quantum memory is a delay line in free-space [136] or optical fibre [137]. This approach is inherently noise free, as it does not involve any atoms or optical fields that could generate noise, and can store broadband signals. However, a simple delay line does not offer the capability of on-demand read out. Instead a memory can be built using a short delay line and a fast switch to allow quasi-on-demand retrieval of the signal. The switch needs to be low loss to enable long storage times, and fast to enable high operating speeds of the memory. This has been demonstrated using a fast Pockels cell with a 4 ns rise time and a round-trip transmission of 98.8%, leading to a memory lifetime of 83 round-trips, or 830 ns [138]. Whilst these results are impressive, it will be challenging to reduce the loss further to enable longer lifetimes, or to scale up the technology to multiple memories using current bulk Pockels cell infrastructure.

An alternative approach is to use light-matter interactions to store optical fields. This can be achieved using single emitters such as a trapped ions [139], an atom in a cavity [140], or colour centres in diamond such as nitrogen-vacancies [141] or silicon-vacancies [142]. However in order to reach a strong coupling regime with a single isolated system it is necessary to use a high finesse cavity [140] or nanophotonic structures [143] to enhance the light-matter interaction, which places a limit on the acceptance bandwidth of the memory and increases the engineering complexity. Instead, one can use an ensemble of atoms to enhance the light-matter coupling without decreasing the bandwidth of signals that can be stored. In an ensemble-based memory, the probability that a photon is absorbed increases with the number of atoms, and the photonic state is mapped into a collective excitation across the ensemble.

Ensemble-based memories have seen a huge amount of research over the past few years, and the different protocols can be separated into two main approaches: optically-controlled memories, and photon-echo based memories. The physical platforms that have been investigated include rare-earth ion-doped solids (REIDs), diamond colour centres, alkali metal vapours and molecules, and most of these platforms can support multiple different memory protocols. Recent results demonstrate memories with efficiencies over 90% [144], storage lifetimes exceeding one second [145], and terahertz storage bandwidths [146], but, to date, there is no single system that has demonstrated all key attributes simultaneously. The remaining challenge is to reproduce these results in a single memory protocol.

1.3.3(a) Photon Echo Memories

Photon-echo memories are based on an ensemble of two-level systems, where the transition has been broadened inhomogeneously by, for example, local strain in a solid state environment, or external electric or magnetic fields. When an inhomogeneously broadened ensemble of atoms absorbs a pulse of light, each atom is excited into a superposition of the ground and excited state, and these coherences each will accumulate phase at a different rate due to the varying detuning, leading to dephasing. Photon echo based memories aim to control the rephasing of these coherences, such that the signal is retrieved from the ensemble in a "photon-echo" [147].

The gradient echo memory (GEM) [148], or the closely related controlled reversible inhomogeneous broadening (CRIB) [149], use an external magnetic field to artificially broaden a narrow absorption line by the Zeeman shift. Once a signal has been absorbed, the broadening causes the atoms to accumulate phase. After a time τ the direction of the field is reversed and the atoms all rephase after time 2τ and the signal is retrieved. GEM has been implemented in an ensemble of rubidium atoms with an efficiency of 87% and a storage time of several μs [150]. In cold atomic ensembles, much longer storage times of 195 μs have been achieved [151]. However, for a given optical depth of the vapour, there is a compromise between memory efficiency and the storage bandwidth, and typically only MHz-bandwidth signals can be stored efficiently, which limits the time-bandwidth product.

A different photon-echo protocol prepares the ensemble in a periodic atomic frequency comb (AFC), where the absorption lines are spaced by a fixed amount δ_c. The periodic absorption feature means that the ensemble rephases periodically every $\tau = 2\pi/\delta_c$. The AFC protocol has been implemented in REIDs [152] and efficiencies of 58% have been demonstrated [153]. How-

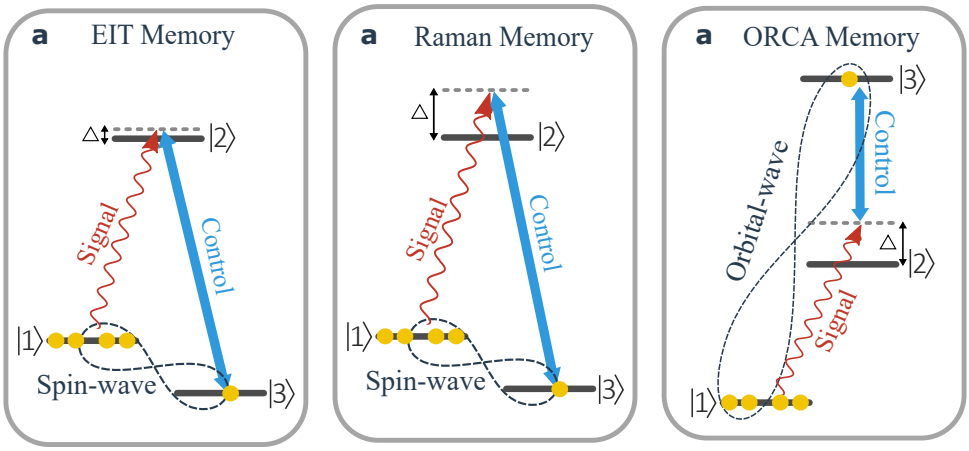


Figure 1.3: Energy level schematics for the EIT (a), Raman (b) and ORCA (c) memory protocols. In EIT and Raman memories, the signal is coherently mapped to a collective excitation between the long-lived spin ground states, or "spin-wave". In the ORCA protocol, the signal is mapped to a coherence between the doubly excited orbital state $|3\rangle$ and the ground state, and this is referred to as an "orbital-wave". More details about the memory protocols are given in the text.

ever, the storage time of the memory is pre-determined by the spacing of the frequency comb, and therefore it does not offer on-demand readout. This can be circumvented by coherently transferring the excitation to a dark "shelf" state for a controllable storage time, known as the full AFC protocol [154], and an efficiency of 12% with on-demand recall has been demonstrated [155]. Storage of polarisation [156] and time-bin qubits [157] at the single photon level have been demonstrated in a full AFC memory, as well as on-demand storage and retrieval of single photons [158].

Photon-echo based memories are inherently multimode due to the broad absorption profile, and storage of multiple temporal [159, 160] and spectral [161, 162] modes have been demonstrated.

1.3.3(b) Optically-Controlled Memories

In an optically-controlled memory, a strong *control* pulse is used to mediate absorption and retrieval of a weak *signal* pulse. The retrieval of the signal is "on-demand", as it occurs on the application of a control pulse. However, the strong control field introduces noise processes which can decrease the fidelity of the memory protocol. Optically-controlled memories are typically based on an ensemble of atoms with three energy levels in a Λ or ladder configuration, as shown in Figure 1.3. The signal and control field are in two-photon resonance, but can be detuned from the excited state by Δ, leading to two main regimes depending on the relative size of the detuning, Δ, and linewidth of the excited state, Γ.

When Δ is small compared to Γ, the control field induces a transparency window in the absorption spectrum of the signal, in a phenomenon known as electromagnetically-induced transparency (EIT). EIT is due to destructive quantum interference between two dressed states which are created by the control field, and allows for the transmission of a signal through the otherwise absorptive medium [163] This reduction in absorption is accompanied by a strong

dispersive feature, and EIT can be used to dramatically decrease the group velocity of the signal [164]. By adiabatically decreasing the control field intensity, the dispersion diverges and the signal light can be slowed to a halt in the medium; mapping it to a collective coherence in the ensemble [165]. By turning the control field back on, the signal field can be re-accelerated and retrieved from the medium.

EIT memories were first proposed by Fleichhauer and Lukin in 2000 [166] and a theoretical model of EIT storage was developed by Gorshkov in 2007 [167]. EIT memories have been demonstrated in warm [168] and cold [165] atomic ensembles and in REIDs [169], and efficiencies of 90.6% have been reached [144]. Storage times of several minutes have been demonstrated in cold atomic ensembles [44], and of one minute in REIDs [170]. Single photon storage has been demonstrated in cold [171] and warm [172] alkali vapour, and polarisation qubits at the single-photon level have been stored with a fidelity above 99% and an efficiency of 85% [144]. The EIT protocol has also been extended to Rydberg state excitations [105, 173], which highlights a route towards single-photon level non-linearities. However, the acceptance bandwidth of EIT protocols is set by the transparency window, which is typically on the MHz scale.

An alternative memory protocol was recently demonstrated, using a phenomenon referred to as the Autler-Townes splitting (ATS) in laser cooled rubidium [174]. ATS is a well-studied phenomenon and is closely related to EIT, but in a strong field regime such that the transparency window broadens, but the quantum interference key to EIT type memories vanishes. ATS was investigated in the context of light-storage for the first time and a memory efficiency of 7% and lifetime of 330 ns was demonstrated for coherent states with a bandwidth of 11 MHz.

To enable more broadband signals to be stored, off-resonant protocols can be used, although the requisite optical depth and control Rabi frequency are significantly higher. In the far-detuned case, when $\Delta \gg \Gamma$, storage and retrieval of optical signals can occur via off-resonant Raman scattering [175, 176]. The bandwidth of the Raman memory is determined by the bandwidth of the control pulse which enables broadband storage, and GHz signals have been stored in warm alkali vapours [177] and molecular hydrogen [178], whilst THz storage has been demonstrated in diamond [146]. By optimising the temporal profile of the control pulse, efficiencies of 82% have been demonstrated in warm rubidium vapour [179], and storage times of 1.5 μs have been shown in warm caesium vapour leading to a time-bandwidth product of $B \sim 2500$ [180]. High fidelity polarisation storage has been performed in warm vapour [181] and in a cold ensemble [182]. Single-photon storage has been demonstrated in room-temperature diamond [146] but the photon number statistics were modified by inherent four-wave mixing type noise, and this same noise process has prevented quantum-level storage in atomic vapour [183].

Recently a similar protocol using off-resonant scattering was proposed, referred to as off-resonant cascaded absorption (ORCA) or fast ladder memory (FLAME). Here the energy level configuration is a ladder rather than Λ configuration (see Figure 1.3(c)), and storage of GHz bandwidth signals has been demonstrated in warm caesium [123] and warm rubidium [184] vapour. This protocol does not suffer from four-wave mixing noise, and storage of single photons with no change in the photon number statistics has been demonstrated ($g^{(2)}_{\text{in}} = g^{(2)}_{\text{out}} = 0.02$ [123]). However the lifetime of the memory was only 5 ns in caesium and 86 ns in rubidium, and hence its applications will be limited to local synchronisation.

The DLCZ protocol introduced in Section 1.3.1(a) can be thought of as a hybrid between a single photon source and an optically-controlled memory. Applying a control pulse drives spontaneous Raman scattering which generates a photon and an atomic excitation in the ensemble. The atomic excitation can be retrieved on demand as a single photon by applying a second control pulse. The DLCZ protocol does not allow storage of input states from other sources, and therefore cannot be used for synchronisation of probabilistic photon sources or logical operations, but it is ideally suited for distributing entanglement over large distances in a quantum repeater scheme. DLCZ memories have been demonstrated in cold rubidium with an efficiency of 76% and a lifetime of 0.22 s [185]. The DLCZ protocol can also be used as a photon source which is bandwidth-compatible with other memory protocols, and storage of DLCZ photons in EIT memories has been demonstrated [144, 171, 172]. The photons emitted from DLCZ sources can be in many spatial modes, and this has been exploited to enable angularly-multiplexed memories [186, 187]. However, we typically require photon sources to emit in a single well-defined spatial mode, and recently a DLCZ protocol has been demonstrated using an array of single atoms trapped along an optical nanofibre which enables on-demand read-out of single photons into the guided mode [188].

1.3.3(c) Summary

Optical quantum memories have been demonstrated in a variety of protocols across many physical platforms. Different realisations of quantum memories have different attributes and therefore are suited to different applications.

In this thesis we focus on Raman memories in warm caesium vapour. Raman memories are broadband and have microsecond lifetimes and therefore are ideally suited to local synchronisation applications. However, the protocol suffers from inherent four-wave mixing noise which has prevented quantum-level demonstrations. In this first half of this thesis I focus on increasing the signal to noise ratio in the Raman memory by increasing the memory efficiency (Chapter 3) and reducing the noise (Chapters 4 and 5). In the final section I focus on the modal properties of the Raman memory and investigate how single mode the memory is and the applications this leads to (Chapter 6). In Chapter 2 I give an in-depth description of the Raman memory, and a detailed outline of this thesis is given in Section 2.3.

Chapter 2:

The Raman Quantum Memory

In this chapter I introduce the Raman memory in more detail and give an outline of previous experimental results. I give a detailed theoretical framework to describe the system, which has been developed by J.N, and explore how to improve upon pre-existing implementations of the memory.

2.1 Raman Memory Protocol

The Raman quantum memory protocol was first proposed in 2000 [175] and is based on an ensemble of atoms each with a Λ-energy level configuration, as shown in Figure 2.1 [176]. The ground states $|1\rangle$ and $|3\rangle$ are long lived, and the system is initialised with all the atoms in the initial state, $|1\rangle$. The signal field is detuned from the $|1\rangle \rightarrow |2\rangle$ atomic transition by Δ and therefore, in the absence of any other fields, undergoes minimal linear absorption and will be transmitted straight through the ensemble. A strong control pulse is applied which is in two-photon resonance with the signal and drives a stimulated, off-resonant two-photon Raman transition from $|1\rangle \rightarrow |3\rangle$. This coherently stores the optical field as a collective excitation of the ground-state coherence, or spin-wave, across the ensemble of atoms. To retrieve the signal, a second control pulse is applied which drives the reverse process and coherently converts the atomic spin-wave back into an optical field.

The Raman memory coherently maps a propagating input optical field to a stationary atomic excitation. The stored mode is a coherence between two long-lived ground states and therefore the storage time of the Raman memory can be long[1]. In the case of a single photon stored in an ideal memory, the spin-wave is a large superposition state where a single atom in the ensemble is in the storage state, $|3\rangle$, and all other atoms are in the initial state, $|1\rangle$. It is not known which atom has changed its internal state, and therefore this is a delocalised excitation across

[1] In practice, the lifetime of the memory is reduced by dephasing mechanisms such as magnetic fields, the motion of the atoms with respect to the phase information of the stored excitation, and the motion of the atoms out of the interaction region.

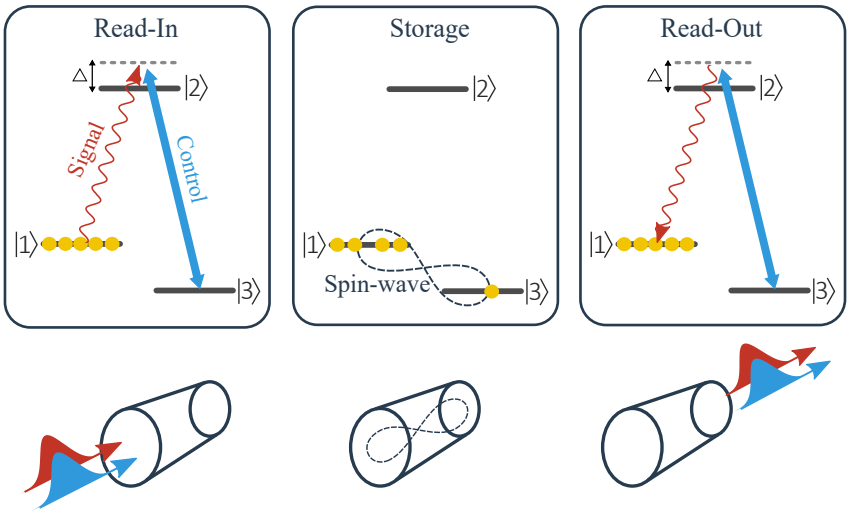


Figure 2.1: The Raman Memory protocol uses a strong control field to drive a Raman transition and store an input signal field as a coherence across an atomic ensemble, or spin-wave. The signal is retrieved from the memory at a later time by applying a read-out control pulse which drives the reverse process and converts the spin-wave back into an optical field.

the entire ensemble. The bandwidth of the memory is determined by the bandwidth of the control pulse, which, in order to address only a single ground state, is upper bounded by the ground state splitting. In typical operational conditions this separation is several GHz, so the Raman memory can store broadband signals with bandwidths $\delta_s \gtrsim 1$ GHz. The large detuning from the excited state reduces noise such as collision-induced fluorescence, and also means that the Doppler broadening of the state is circumvented and hence we can operate the memory at ambient temperatures[2]. However, at large detunings the light-matter coupling becomes weak, and we must interact with a large ensemble of atoms to achieve a collective enhancement of the coupling strength.

The Raman memory has been demonstrated experimentally in warm caesium vapour using the Λ configuration of the D_2 atomic transition. More details about the experimental infrastructure for implementing the Raman memory are given in Chapter 3, but the key results are summarised below. In 2010, a gigahertz-bandwidth signal was stored in the Raman memory with a total efficiency of storage and retrieval of $\eta_{\text{mem}} = 15\%$ for a storage time of $\tau_{\text{mem}} = 12.5$ ns [177]. The retrieved state was interfered with a delayed copy of the input state, and interference fringes with a maximum visibility of $(86 \pm 5)\%$ were observed, demonstrating the coherent nature of the memory interaction. Later, storage of weak coherent states with an average photon number of $N_{\text{in}} \sim 1.6$ was demonstrated with an efficiency of $\eta_{\text{mem}} = 30\%$ and memory lifetime of $\tau_{\text{mem}} = 1.5\ \mu$s [180], giving a large time-bandwidth product $B \sim 2500$. The Raman memory protocol in alkali vapours has therefore been demonstrated to combine high-bandwidth storage, moderate efficiencies and lifetimes, and technical simplicity due to room-temperature operation, making them a promising candidate for broadband quantum memories.

However, in 2015 the Raman memory was interfaced with a single photon source, and

[2]We often refer to the Raman memory as a "room temperature" protocol as it does not require cryogenics or laser cooling techniques. Typical operational temperatures are around 70 - 80°C, or a very warm room...

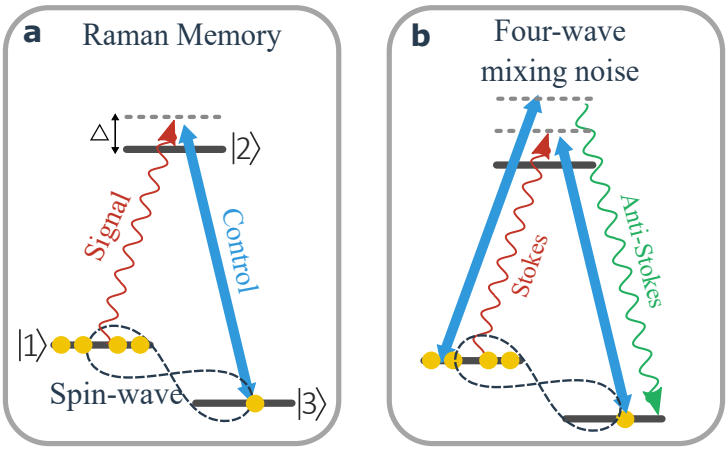


Figure 2.2: (a) A schematic of the Raman memory protocol. (b) Four-wave mixing noise arises when the control field couples to the populated state $|1\rangle$ and drives anti-Stokes scattering.

heralded single photons with $g^{(2)}_{\text{in}} = 0.016 \pm 0.004$ were stored in the memory and retrieved with $g^{(2)}_{\text{out}} = 1.59 \pm 0.03$ [183], which is significantly above the non-classical threshold of $g^{(2)} = 1$. The increase in $g^{(2)}$ was attributed to four-wave mixing noise which coherently adds to the retrieved state and destroys the non-classicality of the photon number statistics. The origin of the four-wave mixing noise is the unwanted coupling of the strong control field to the populated ground state, $|1\rangle$, which drives spontaneous anti-Stokes scattering, as shown in Figure 2.2(b). This creates a *noisy* spin-wave which has significant overlap with the *memory* spin-wave, and is efficiently read out in the same temporal and spectral mode as the signal field. This adds thermal noise to the retrieved photonic state, preventing storage at the single-photon level. Suppressing four-wave mixing noise has therefore been identified as a key challenge to enable quantum storage and retrieval in the Raman memory.

The primary goal of this thesis is to increase the signal to noise ratio in the Raman memory, by both maximising the efficiency of the protocol and suppressing four-wave mixing noise. In the following section we will derive the equations of motion that describe the Raman memory protocol, as this framework allows us to investigate different methods to increase the signal to noise ratio.

2.2 Theory

In this Section I present a more formal description of the Raman memory by re-deriving the equations of motion that govern the system. This derivation closely follows those presented in [189, 190], and captures the desired memory interaction as well as the four-wave mixing noise process.

2.2.1 Optical Bloch Equations

To derive the equations of motion for the Raman memory we consider an ensemble of atoms with a Λ energy level configuration, as shown in Figure 2.1(a). The atomic eigenstates are labelled by $|i\rangle$ for $i = 1, 2, 3$, and the ensemble is initialised into long-lived ground state $|1\rangle$.

The input optical signal field, $\boldsymbol{E}_s$, is stored as an atomic coherence between states $|1\rangle$ and $|3\rangle$ via an off-resonant Raman transition driven by a strong control field $\boldsymbol{E}_c$.

The total Hamiltonian for this system can be written as:

$$\hat{H} = \hat{H}_A + \hat{H}_L + \hat{H}_{\text{int}}, \tag{2.1}$$

where $\hat{H}_A$ is the Hamiltonian for an atom in the ensemble, $\hat{H}_L$ is the Hamiltonian for the light fields, and the interaction Hamiltonian, $\hat{H}_{\text{int}}$, describes the electric dipole light-matter interaction. The atomic Hamiltonian can be written in terms of the energy eigenstates as:

$$\hat{H}_A = \sum_{j=1}^{3} \omega_j \, |j\rangle\langle j| = \sum_{j=1}^{3} \omega_j \, \hat{\sigma}_{jj} \tag{2.2}$$

where we have dropped $\hbar$ for simplicity, and $\hat{\sigma}_{jk} = |j\rangle\langle k|$ are flip operators satisfying

$$\begin{aligned} \hat{\sigma}_{ij}\hat{\sigma}_{kl} &= \hat{\sigma}_{il}\,\delta_{jk} \\ \hat{\sigma}_{jk}^{\dagger} &= \hat{\sigma}_{kj}. \end{aligned} \tag{2.3}$$

The diagonal operators $\hat{\sigma}_{jj}$ describe the atomic populations in each state, and the off-diagonal $\hat{\sigma}_{j,k\neq j}$ describe the atomic coherences between them.

To consider the light-field Hamiltonian, $\hat{H}_L$, we write expressions for the optical fields present. The control pulse is sufficiently intense that it can be represented as a classical field:

$$\boldsymbol{E}_c(t, z, \boldsymbol{\rho}) = \boldsymbol{v}_c E_c(t, z, \boldsymbol{\rho}) \mathrm{e}^{\mathrm{i}\omega_c(t-z/c)} + \text{c.c.} \tag{2.4}$$

where $\boldsymbol{v}_c$ is the polarisation vector of the field, E_c is the slowly varying amplitude envelope, and ω_c is the central angular frequency. Here z is the direction of propagation of the field, and $\boldsymbol{\rho}$ is the transverse position vector. The signal field needs to be treated quantum mechanically as we wish to store non-classical states of light, and hence is expressed as

$$\hat{\boldsymbol{E}}_s(t, z, \boldsymbol{\rho}) = \mathrm{i}\boldsymbol{v}_s \int \mathrm{d}\omega \, g(\omega) \hat{a}(\omega, t, \boldsymbol{\rho}) \mathrm{e}^{-\mathrm{i}\omega z/c} + \text{h.c.} \tag{2.5}$$

where $\hat{a}(\omega, t, \boldsymbol{\rho})$ is the Bosonic creation operator for a photon with frequency ω at time t and position $\boldsymbol{\rho}$, and $g(\omega) = \sqrt{\hbar\omega/4\pi\epsilon_0 cA}$ is the mode amplitude at frequency ω for a beam of transverse mode area A. If the bandwidth of the signal field, δ_s, is much less than the central frequency, ω_s, then we can neglect the ω-dependence of the amplitude and this reduces to:

$$\hat{\boldsymbol{E}}_s = \mathrm{i}\boldsymbol{v}_s g_s \hat{S}(t, z, \boldsymbol{\rho}) \mathrm{e}^{\mathrm{i}\omega_s(t-z/c)} + \text{h.c.} \tag{2.6}$$

where $g_s = \sqrt{2\pi} g(\omega_s)$ and

$$\hat{S}(t, z, \boldsymbol{\rho}) = \mathrm{e}^{-\mathrm{i}\omega(t-z/c)} \times \frac{1}{\sqrt{2\pi}} \int \mathrm{d}\omega \, \hat{a}(\omega, t, \boldsymbol{\rho}) \mathrm{e}^{-\mathrm{i}\omega z/c} \tag{2.7}$$

is the slowly-varying mode amplitude of the signal.

The light-matter Hamiltonian is given by the electric dipole interaction

$$\hat{H}_{\text{int}} = -\boldsymbol{E} \cdot \hat{\boldsymbol{d}} \tag{2.8}$$

where $\hat{\boldsymbol{d}}$ is the dipole operator. The dipole moment can be written in terms of its matrix elements as

$$\hat{\boldsymbol{d}} = \sum_{j,k} \langle j| \hat{\boldsymbol{d}} |k\rangle \, \hat{\sigma}_{jk} = \sum_{j,k} \boldsymbol{d}_{jk} \hat{\sigma}_{jk}. \tag{2.9}$$

The diagonal elements $\boldsymbol{d}_{jj}$ are zero as the electric dipole operator has odd parity, and $\boldsymbol{d}_{13} = 0$ as the transition $|1\rangle \leftrightarrow |3\rangle$ is dipole-forbidden, and hence Equation 2.9 reduces to

$$\hat{\boldsymbol{d}} = \boldsymbol{d}_{12}\hat{\sigma}_{12} + \boldsymbol{d}_{23}\hat{\sigma}_{23} + \text{h.c.} \tag{2.10}$$

To describe the evolution of the system we consider the Heisenberg equations of motion

$$\partial_t \hat{\sigma}_{jk} = i[\hat{\sigma}_{jk}, \hat{H}] = i[\hat{\sigma}_{jk}, \hat{H}_A + \hat{H}_{\text{int}}] \tag{2.11}$$

which gives the following coupled equations:

$$\begin{aligned}
\partial_t \hat{\sigma}_{11} &= -\mathrm{i}\boldsymbol{E} \cdot [\boldsymbol{d}_{12}\hat{\sigma}_{12} + \text{h.c.}] && (2.12)\\
\partial_t \hat{\sigma}_{33} &= +\mathrm{i}\boldsymbol{E} \cdot [\boldsymbol{d}_{23}\hat{\sigma}_{23} + \text{h.c.}] && (2.13)\\
\partial_t \hat{\sigma}_{12} &= \mathrm{i}\omega_{21}\hat{\sigma}_{12} - \mathrm{i}\boldsymbol{E} \cdot [\boldsymbol{d}_{12}^*(\hat{\sigma}_{11} - \hat{\sigma}_{22}) + \boldsymbol{d}_{23}\hat{\sigma}_{13}] && (2.14)\\
\partial_t \hat{\sigma}_{13} &= \mathrm{i}\omega_{31}\hat{\sigma}_{13} - \mathrm{i}\boldsymbol{E} \cdot [\boldsymbol{d}_{23}^*\hat{\sigma}_{12} - \boldsymbol{d}_{12}^*\hat{\sigma}_{23}] && (2.15)\\
\partial_t \hat{\sigma}_{23} &= \mathrm{i}\omega_{32}\hat{\sigma}_{23} - \mathrm{i}\boldsymbol{E} \cdot [\boldsymbol{d}_{23}^*(\hat{\sigma}_{22} - \hat{\sigma}_{33}) - \boldsymbol{d}_{12}\hat{\sigma}_{13}] && (2.16)
\end{aligned}$$

where $\omega_{jk} = \omega_j - \omega_k$. We have used the fact that the total population is conserved $\left(\sum_j \hat{\sigma}_{jj} = \mathbb{1}\right)$ to eliminate the equation for $\partial_t \hat{\sigma}_{22}$. We note that the origin of four-wave mixing noise is spontaneous anti-Stokes scattering arising from the unwanted coupling of the control field to the $|1\rangle \leftrightarrow |2\rangle$ transition via $\boldsymbol{d}_{12} \cdot \boldsymbol{v}_c$. This interaction is detuned by $\Delta_+ = \Delta + \Delta_{\text{hf}}$, and we hence also need to consider a weak anti-Stokes field at the frequency $\omega_a = \omega_c + \omega_{13}$:

$$\hat{\boldsymbol{E}}_a = \mathrm{i}\boldsymbol{v}_a g_a \hat{A}(t, z, \boldsymbol{\rho}) \mathrm{e}^{\mathrm{i}\omega_a(t - z/c)} + \text{h.c.} \tag{2.17}$$

To gain physical insight from these equations we make a series of approximations.

1. **Linear approximation (weak signal):**
   In the limit of weak signals, where the number of photons stored is much less than the total number of atoms, we can assume that the atomic populations are constant in time. We therefore set $\partial_t \hat{\sigma}_{jj} = 0$, and replace the population operators with their expectation values $\hat{\sigma}_{11} \to (1 - \alpha) \approx 1$, $\hat{\sigma}_{22} \to 0$, and $\hat{\sigma}_{33} \to \alpha \approx 0$. Here $(1 - \alpha)$ is the average proportion of atoms in the initial state $|1\rangle$, or the optical pumping efficiency, and $\alpha \neq 0$ accounts for imperfect initialisation of the atomic ensemble.

2. **Rotating wave approximation:**
   We can transform to a rotating frame by defining $\tilde{\sigma}_{jk} = \hat{\sigma}_{jk} e^{i\omega_{jk}\tau}$ where $\tau = t - z/c$

is the retarded time. Provided that the detuning, Δ, is much less than the control and signal frequencies, $\omega_{s,c}$, then we can use the rotating wave approximation and neglect fast-oscillating terms.

This transformation gives:

$$\begin{aligned}
\partial_\tau \tilde{\sigma}_{12} &= -\mathrm{i}\boldsymbol{E}\cdot\left((1-\alpha)\boldsymbol{d}_{12}^{*}\mathrm{e}^{-\mathrm{i}\omega_{21}\tau}+\boldsymbol{d}_{23}\tilde{\sigma}_{13}\mathrm{e}^{-\mathrm{i}\omega_{23}\tau}\right) && (2.18)\\
\partial_\tau \tilde{\sigma}_{13} &= -\mathrm{i}\boldsymbol{E}\cdot\left(\boldsymbol{d}_{23}^{*}\tilde{\sigma}_{12}\mathrm{e}^{+\mathrm{i}\omega_{23}\tau}-\boldsymbol{d}_{12}^{*}\tilde{\sigma}_{23}\mathrm{e}^{-\mathrm{i}\omega_{21}\tau}\right) && (2.19)\\
\partial_\tau \tilde{\sigma}_{23} &= +\mathrm{i}\boldsymbol{E}\cdot\left(\alpha\boldsymbol{d}_{23}^{*}\mathrm{e}^{+\mathrm{i}\omega_{23}\tau}+\boldsymbol{d}_{12}\tilde{\sigma}_{13}\mathrm{e}^{+\mathrm{i}\omega_{21}\tau}\right). && (2.20)
\end{aligned}$$

We write the electric field as

$$\begin{aligned}
\boldsymbol{E} &= \boldsymbol{E}_c+\boldsymbol{E}_s+\boldsymbol{E}_a\\
&= \boldsymbol{c}\mathrm{e}^{\mathrm{i}\omega_c\tau}+\boldsymbol{s}\mathrm{e}^{\mathrm{i}\omega_s\tau}+\boldsymbol{a}\mathrm{e}^{\mathrm{i}\omega_a\tau}+\mathrm{h.c.} && (2.21)
\end{aligned}$$

where, for simplicity, we have defined $\boldsymbol{c}=E_c(t,z,\boldsymbol{\rho})\boldsymbol{v}_c$, $\boldsymbol{s}=\mathrm{i}g_s\hat{S}(t,z,\boldsymbol{\rho})\boldsymbol{v}_s$ and $\boldsymbol{a}=\mathrm{i}g_a\hat{A}(t,z,\boldsymbol{\rho})\boldsymbol{v}_a$. We substitute Equation 2.21 into Equations 2.18 – 2.20, and only keep terms with frequencies $\pm(\omega_{ij}-\omega_{c,s,a})$, since terms oscillating at $\pm(\omega_{ij}+\omega_{c,s,a})$ will average to zero (from the rotating wave approximation). We define the detuning from resonance as $\Delta=\omega_s-\omega_{21}$, such that Δ is positive for a blue detuned signal. The ground state splitting is given by $\Delta_{\mathrm{hf}}=\omega_{13}$. This gives:

$$\begin{aligned}
\partial_\tau \tilde{\sigma}_{12} &= -\mathrm{i}(1-\alpha)\boldsymbol{d}_{12}^{*}\cdot\left[\boldsymbol{s}\mathrm{e}^{\mathrm{i}\Delta\tau}+\boldsymbol{c}\mathrm{e}^{\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}+\boldsymbol{a}\mathrm{e}^{\mathrm{i}(\Delta+2\Delta_{\mathrm{hf}})\tau}\right]\\
&\quad -\mathrm{i}\tilde{\sigma}_{13}\boldsymbol{d}_{23}\cdot\left[\boldsymbol{s}\mathrm{e}^{\mathrm{i}(\Delta-\Delta_{\mathrm{hf}})\tau}+\boldsymbol{c}\mathrm{e}^{\mathrm{i}\Delta\tau}+\boldsymbol{a}\mathrm{e}^{\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}\right] && (2.22)\\
\partial_\tau \tilde{\sigma}_{13} &= -\mathrm{i}\tilde{\sigma}_{12}\boldsymbol{d}_{23}^{*}\cdot\left[\boldsymbol{s}^{\dagger}\mathrm{e}^{-\mathrm{i}(\Delta-\Delta_{\mathrm{hf}})\tau}+\boldsymbol{c}^{*}\mathrm{e}^{-\mathrm{i}\Delta\tau}+\boldsymbol{a}^{\dagger}\mathrm{e}^{-\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}\right]\\
&\quad +\mathrm{i}\tilde{\sigma}_{23}\boldsymbol{d}_{12}^{*}\cdot\left[\boldsymbol{s}\mathrm{e}^{\mathrm{i}\Delta\tau}+\boldsymbol{c}\mathrm{e}^{\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}+\boldsymbol{a}\mathrm{e}^{\mathrm{i}(\Delta+2\Delta_{\mathrm{hf}})\tau}\right] && (2.23)\\
\partial_\tau \tilde{\sigma}_{23} &= +\mathrm{i}\alpha\boldsymbol{d}_{23}^{*}\cdot\left[\boldsymbol{s}^{\dagger}\mathrm{e}^{-\mathrm{i}(\Delta-\Delta_{\mathrm{hf}})\tau}+\boldsymbol{c}^{*}\mathrm{e}^{-\mathrm{i}\Delta\tau}+\boldsymbol{a}^{\dagger}\mathrm{e}^{-\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}\right]\\
&\quad +\mathrm{i}\tilde{\sigma}_{13}\boldsymbol{d}_{12}\cdot\left[\boldsymbol{s}^{\dagger}\mathrm{e}^{-\mathrm{i}\Delta\tau}+\boldsymbol{c}^{*}\mathrm{e}^{-\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}+\boldsymbol{a}^{\dagger}\mathrm{e}^{-\mathrm{i}(\Delta+2\Delta_{\mathrm{hf}})\tau}\right]. && (2.24)
\end{aligned}$$

We next consider the propagation of the weak signal field which is described by Maxwell's equations:

$$\left(\nabla^2-\frac{1}{c^2}\partial_t^2\right)\boldsymbol{E}_s=\mu_0\partial_t^2\boldsymbol{P}_s \tag{2.25}$$

where $\boldsymbol{P}_s$ is the positive frequency component of the dipole moment per unit volume, or atomic polarisation, which oscillates at the signal frequency ω_s and acts as a source for the signal field. We rewrite $\boldsymbol{P}_s$ in terms of the slowly varying polarisation, $\tilde{\boldsymbol{P}}_s$ and the fast oscillating carrier wave: $\boldsymbol{P}_s=\tilde{\boldsymbol{P}}_s\mathrm{e}^{\mathrm{i}\omega_s\tau}$, which gives:

$$\left(\nabla^2-\frac{1}{c^2}\partial_t^2\right)\left[\boldsymbol{s}\mathrm{e}^{\mathrm{i}\omega_s\tau}\right]=\mu_0\partial_t^2\tilde{\boldsymbol{P}}_s\mathrm{e}^{\mathrm{i}\omega_s\tau}. \tag{2.26}$$

We split the spatial derivative, ∇^2, into a transverse component, $\nabla_\perp^2 = \partial_x^2 + \partial_y^2$, and a component along the direction of propagation, ∂_z^2, and use the chain rule of both sides of the above equation:

$$\begin{aligned}\left(\nabla^2 - \frac{1}{c^2}\partial_t^2\right)\left[\boldsymbol{s}\mathrm{e}^{\mathrm{i}\omega_s\tau}\right] &= \mathrm{i}g_s\boldsymbol{v}_s\mathrm{e}^{\mathrm{i}\omega_s\tau}\left\{\nabla_\perp^2 + \left(\partial_z^2 - \frac{1}{c^2}\partial_t^2\right) - \frac{2\mathrm{i}\omega_s}{c}\left(\partial_z + \frac{1}{c}\partial_t\right)\right\}S \\ &\approx \mathrm{i}g_s\boldsymbol{v}_s\mathrm{e}^{\mathrm{i}\omega_s\tau}\left\{\nabla_\perp^2 - 2\mathrm{i}k_s\left(\partial_z + \frac{1}{c}\partial_t\right)\right\}S \end{aligned} \tag{2.27}$$

and

$$\begin{aligned}\partial_t^2\left[\tilde{\boldsymbol{P}}_s\mathrm{e}^{\mathrm{i}\omega_s\tau}\right] &= \mathrm{e}^{\mathrm{i}\omega_s\tau}\left\{\partial_t^2 + 2\mathrm{i}\omega_s\partial_t - \omega_s^2\right\}\tilde{\boldsymbol{P}}_s \\ &\approx -\omega_s^2\mathrm{e}^{\mathrm{i}\omega_s\tau}\tilde{\boldsymbol{P}}_s. \end{aligned} \tag{2.28}$$

We have used the slowly varying envelope approximation to neglect the second order derivatives, and let $k_s = \omega_s/c$. Putting this together gives:

$$\left\{\frac{\mathrm{i}}{2k_s}\nabla_\perp^2 + \left(\partial_z + \frac{1}{c}\partial_t\right)\right\}S = -\frac{\mu_0\omega_s^2}{2g_sk_s}\boldsymbol{v}_s^*\cdot\tilde{\boldsymbol{P}}_s. \tag{2.29}$$

To proceed we consider the polarisation density in terms of the atomic dipole operator:

$$\begin{aligned}\tilde{\boldsymbol{P}}_s &= \mathrm{e}^{-\mathrm{i}\omega_s\tau}\frac{1}{\delta V}\sum_{\beta(\boldsymbol{r})}\boldsymbol{d}^{(\beta)} \\ &= \mathrm{e}^{-\mathrm{i}\omega_s\tau}\frac{1}{\delta V}\sum_{\beta(\boldsymbol{r})}\left[\boldsymbol{d}_{12}\hat{\sigma}_{12}^{(\beta)} + \boldsymbol{d}_{32}\hat{\sigma}_{32}^{(\beta)}\right] \end{aligned} \tag{2.30}$$

where β labels all the atoms in a small volume δV at position $\boldsymbol{r}$. We define the collective polarisation operator driven at the signal frequency as:

$$\begin{aligned}\hat{P} &= \frac{1}{\sqrt{n}\delta V}\sum_{\beta(\boldsymbol{r})}\hat{\sigma}_{12}\mathrm{e}^{-\mathrm{i}\omega_s\tau}, \\ &= \frac{1}{\sqrt{n}\delta V}\sum_{\beta(\boldsymbol{r})}\tilde{\sigma}_{12}\mathrm{e}^{-\mathrm{i}\Delta\tau}, \end{aligned} \tag{2.31}$$

where n is the number density of atoms. Similarly we define a polarisation operator for the $\hat{\sigma}_{32}$ coherence which is a source term for the anti-Stokes field and oscillates at ω_a:

$$\begin{aligned}\hat{R} &= \frac{1}{\sqrt{n}\delta V}\sum_{\beta(\boldsymbol{r})}\hat{\sigma}_{32}\mathrm{e}^{-\mathrm{i}\omega_a\tau}, \\ &= \frac{1}{\sqrt{n}\delta V}\sum_{\beta(\boldsymbol{r})}\tilde{\sigma}_{32}\mathrm{e}^{-\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}, \end{aligned} \tag{2.32}$$

and for the ground state coherence

$$\hat{B} = \frac{1}{\sqrt{n}\delta V} \sum_{\beta(\boldsymbol{r})} \tilde{\sigma}_{13}. \tag{2.33}$$

It can be shown that as $\delta V \to 0$ (the continuum approximation) P, R and B satisfy bosonic commutation relations, and therefore we identify them as annihilation operators for the atomic polarisations and spin-wave coherence respectively [189, 191]. If we substitute Equations 2.31 and 2.32 into Equation 2.30 we get:

$$\tilde{\boldsymbol{P}}_s = \sqrt{n}\left(\boldsymbol{d}_{12}P + \mathrm{e}^{2\mathrm{i}\Delta_{\mathrm{hf}}\tau}\boldsymbol{d}_{32}R\right), \tag{2.34}$$

and then Equation 2.29 gives

$$\left\{\frac{\mathrm{i}}{2k_s}\nabla_\perp^2 + \left(\partial_z + \frac{1}{c}\partial_t\right)\right\} S = -\frac{\sqrt{n}\mu_0\omega_s^2}{2g_s k_s}\left(\boldsymbol{v}_s^* \cdot \boldsymbol{d}_{12}P + \boldsymbol{v}_s^* \cdot \boldsymbol{d}_{32}\mathrm{e}^{2\mathrm{i}\Delta_{\mathrm{hf}}\tau}R\right). \tag{2.35}$$

We make the paraxial approximation by assuming that the beam has negligible divergence in the transverse direction and thereby drop the transverse Laplacian. We can also derive a similar expression for the anti-Stokes field which gives:

$$\left(\partial_z + \frac{1}{c}\partial_t\right) S = -\left(\chi_s^{*(21)}P + \chi_s^{*(23)}\mathrm{e}^{2\mathrm{i}\Delta_{\mathrm{hf}}\tau}R\right) \tag{2.36}$$

$$\left(\partial_z + \frac{1}{c}\partial_t\right) A = -\left(\chi_a^{*(23)}R + \chi_a^{*(21)}\mathrm{e}^{-2\mathrm{i}\Delta_{\mathrm{hf}}\tau}P\right) \tag{2.37}$$

where we have defined

$$\chi_{s,a}^{(ij)} = \frac{\boldsymbol{d}_{ij} \cdot \boldsymbol{v}_{s,a}}{\hbar}\sqrt{n}g_{s,a}. \tag{2.38}$$

To find expressions for P and R we substitute their definitions (Equations 2.31 and 2.32) into Equations 2.22 – 2.24, yielding:

$$\begin{aligned}\partial_t P &= -\mathrm{i}\Delta P + \frac{1}{\sqrt{n}\delta V}\sum_{\beta(\boldsymbol{r})} \mathrm{e}^{-\mathrm{i}\Delta\tau}\partial_t\tilde{\sigma}_{12} \\ &= -\mathrm{i}\Delta P - \mathrm{i}(1-\alpha)\frac{\sqrt{n}}{\hbar}\boldsymbol{d}_{12}^*\cdot\left(\boldsymbol{s}+\boldsymbol{c}\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau}+\boldsymbol{a}\mathrm{e}^{2\mathrm{i}\Delta_{\mathrm{hf}}\tau}\right) - \mathrm{i}\frac{B}{\hbar}\boldsymbol{d}_{23}\cdot\left(\boldsymbol{s}\mathrm{e}^{-\mathrm{i}\Delta_{\mathrm{hf}}\tau}+\boldsymbol{c}+\boldsymbol{a}\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau}\right)\end{aligned} \tag{2.39}$$

$$\begin{aligned}\partial_t R &= -\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})R + \frac{1}{\sqrt{n}\delta V}\sum_{\beta(\boldsymbol{r})} \mathrm{e}^{-\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})\tau}(\partial_t\tilde{\sigma}_{23})^\dagger \\ &= -\mathrm{i}(\Delta+\Delta_{\mathrm{hf}})R - \mathrm{i}\alpha\frac{\sqrt{n}}{\hbar}\boldsymbol{d}_{23}\cdot\left(\boldsymbol{s}\mathrm{e}^{-2\mathrm{i}\Delta_{\mathrm{hf}}\tau}+\boldsymbol{c}\mathrm{e}^{-\mathrm{i}\Delta_{\mathrm{hf}}\tau}+\boldsymbol{a}\right) - \mathrm{i}\frac{B^\dagger}{\hbar}\boldsymbol{d}_{12}^*\cdot\left(\boldsymbol{s}\mathrm{e}^{-\mathrm{i}\Delta_{\mathrm{hf}}\tau}+\boldsymbol{c}+\boldsymbol{a}\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau}\right)\end{aligned} \tag{2.40}$$

$$\begin{aligned}\partial_t B &= \frac{1}{\sqrt{n}\delta V}\sum_{\beta(\boldsymbol{r})}\partial_t\tilde{\sigma}_{13} \\ &= -\mathrm{i}\frac{P}{\hbar}\boldsymbol{d}_{23}^*\cdot\left(\boldsymbol{s}^\dagger\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau}+\boldsymbol{c}^*+\boldsymbol{a}^\dagger\mathrm{e}^{-\mathrm{i}\Delta_{\mathrm{hf}}\tau}\right) + \mathrm{i}\frac{R^\dagger}{\hbar}\boldsymbol{d}_{12}^*\cdot\left(\boldsymbol{s}\mathrm{e}^{-\mathrm{i}\Delta_{\mathrm{hf}}\tau}+\boldsymbol{c}+\boldsymbol{a}\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau}\right)\end{aligned} \tag{2.41}$$

where we have used the fact that, in the continuum limit ($\delta V \to 0$), the sum over the ensemble gives a factor of $n/\hbar$. To account for decoherence we phenomenologically add decay terms to the expressions for P and R. The excited state population decays at a rate of 2γ and so the atomic coherences decay as γ, and we therefore introduce a term $-\gamma P$ to the equation for $\partial_t P$, and equivalently for $\partial_t R$. We assume that the spin-wave does not decay over the timescale of storage of the memory, and thus do not introduce a decoherence term for B. In addition to decay terms, we should also add Langevin noise operators which describe fluctuations and ensure that the bosonic commutation relations of P, R and B still hold. However, these operators will only appear in normally ordered expectation values in expressions for the memory efficiency and noise, and since the operators are initially in the vacuum state these expectation values are zero and we can neglect such terms.

To solve these equations we assume that the coherences P, R and B are small, as are the signal and anti-Stokes fields S and A. We therefore neglect any second order terms in $\{P, R, B, S, A\}$, giving:

$$(\partial_t+\Gamma_s)P = (1-\alpha)\left(\chi_s^{(21)}S - \mathrm{i}\Omega_{21}\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau} + \chi_a^{(21)}\mathrm{e}^{2\mathrm{i}\Delta_{\mathrm{hf}}\tau}A\right) - \mathrm{i}\Omega_{23}B \tag{2.42}$$

$$(\partial_t+\Gamma_a)R = \alpha\left(\chi_s^{(23)}\mathrm{e}^{-2\mathrm{i}\Delta_{\mathrm{hf}}\tau}S - \mathrm{i}\Omega_{23}\mathrm{e}^{-\mathrm{i}\Delta_{\mathrm{hf}}\tau} + \chi_a^{(23)}A\right) - \mathrm{i}\Omega_{21}B^\dagger \tag{2.43}$$

where $\Gamma_s = \gamma + \mathrm{i}\Delta$ ($\Gamma_a = \gamma + \mathrm{i}(\Delta+\Delta_{\mathrm{hf}})$) is the complex decay rate for the signal (anti-Stokes) field, and $\Omega_{ij} = \boldsymbol{d}_{ij}\cdot\boldsymbol{v}_c E_c/\hbar$ is the Rabi frequency of the control field.

We also assume that the only time dependence in Equations 2.42 – 2.43 is explicit in the exponentials, as the system is adiabatically following the applied fields. We can hence solve these equations to find P by looking for a particular solution of the form $c_0 + c_1\mathrm{e}^{-\Gamma t}$, and a complementary solution of the form $c_2\,\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau} + c_3\,\mathrm{e}^{2\mathrm{i}\Delta_{\mathrm{hf}}\tau}$ for some coefficients $c_{0,1,2,3}$. By substituting in a solution of this form we can find the coefficients and hence find P (and similarly for R):

$$P \approx (1-\alpha)\left(\frac{\chi_s^{(21)}}{\Gamma_s}S - \mathrm{i}\frac{\Omega_{21}\mathrm{e}^{\mathrm{i}\Delta_{\mathrm{hf}}\tau}}{\Gamma_a} + \frac{\chi_a^{(21)}\mathrm{e}^{2\mathrm{i}\Delta_{\mathrm{hf}}\tau}}{\Gamma_a+\mathrm{i}\Delta_{\mathrm{hf}}}A\right) - \mathrm{i}\frac{\Omega_{23}}{\Gamma_s}B \tag{2.44}$$

$$R \approx \alpha\left(\frac{\chi_s^{(23)}\mathrm{e}^{-2\mathrm{i}\Delta_{\mathrm{hf}}\tau}}{\Gamma_s-\mathrm{i}\Delta_{\mathrm{hf}}}S - \mathrm{i}\frac{\Omega_{23}\mathrm{e}^{-\mathrm{i}\Delta_{\mathrm{hf}}\tau}}{\Gamma_s} + \frac{\chi_a^{(23)}}{\Gamma_a}A\right) - \mathrm{i}\frac{\Omega_{21}}{\Gamma_a}B^\dagger. \tag{2.45}$$

Substituting these into Equation 2.41 and keeping only stationary terms (the rotating wave approximation again) as well as neglecting higher order terms in $\{P, R, B, S, A\}$ gives:

$$\begin{aligned}\partial_t B &= -\mathrm{i}\Omega_{23}^*\chi_s^{(21)}\left[\frac{(1-\alpha)}{\Gamma_s}+\frac{\alpha}{\Gamma_s^*}\right]S + \mathrm{i}\Omega_{12}^*\chi_a^{*(23)}\left[\frac{(1-\alpha)}{\Gamma_a}+\frac{\alpha}{\Gamma_a^*}\right]A \\ &\quad -\left[\frac{|\Omega_{12}|^2}{\Gamma_a^*}+\frac{|\Omega_{23}|^2}{\Gamma_s}\right]B.\end{aligned} \tag{2.46}$$

Finally we substitute the expressions for P and R from Equations 2.44 and 2.45 into the wave equations for the signal and anti-Stokes fields (Equations 2.36 and 2.37), and arrive at the equations of motion for the Raman memory:

$$\begin{aligned}\partial_t B &= -\mathrm{i}\Omega_{23}^*\sqrt{\frac{d\gamma}{L}}\left[\frac{(1-\alpha)}{\Gamma_s}+\frac{\alpha}{\Gamma_s^*}\right]S + \mathrm{i}\Omega_{12}^*\sqrt{\frac{d\gamma}{L}}\left[\frac{(1-\alpha)}{\Gamma_a}+\frac{\alpha}{\Gamma_a^*}\right]A^\dagger \\ &\quad -\left[\frac{|\Omega_{12}|^2}{\Gamma_a^*}+\frac{|\Omega_{23}|^2}{\Gamma_s}\right]B\end{aligned} \tag{2.47}$$

$$\left(\partial_z+\frac{1}{c}\partial_t\right)S = \mathrm{i}\sqrt{\frac{d\gamma}{L}}\frac{\Omega_{23}}{\Gamma_s}B - \frac{d\gamma}{L}\left[\frac{(1-\alpha)}{\Gamma_s}+\frac{\alpha}{\Gamma_s^*+\mathrm{i}\Delta_{\mathrm{hf}}}\right]S \tag{2.48}$$

$$\left(\partial_z+\frac{1}{c}\partial_t\right)A = \mathrm{i}\sqrt{\frac{d\gamma}{L}}\frac{\Omega_{12}}{\Gamma_a}B^\dagger - \frac{d\gamma}{L}\left[\frac{(1-\alpha)}{\Gamma_a+\mathrm{i}\Delta_{\mathrm{hf}}}+\frac{\alpha}{\Gamma_a^*}\right]A \tag{2.49}$$

We have defined the optical depth as $d = \frac{|\chi|^2 L}{\gamma}$, where L is the length of the atomic medium. Experimentally d is a measure of the proportion of monochromatic light that will be absorbed through an atomic medium: $I_{\mathrm{out}} = I_{\mathrm{in}}\mathrm{e}^{-2d}$. We shall see below that the optimal efficiency of the Raman memory depends only on the optical depth and therefore d is a key parameter.

These equations describe both the Raman memory interaction which, as we see more explicitly below, is equivalent to a beam splitter Hamiltonian between modes S and B, as well as four-wave mixing noise via the two-mode squeezing interaction between A and B. The final term in Equation 2.47 accounts for the AC Stark shift which is proportional to $|\Omega|^2/\Gamma$, and the final terms in Equations 2.48 and 2.49 account for absorption and dispersion of the signal and anti-Stokes fields.

2.2.2 Modal Properties of the Raman Memory

In the ideal case we can set the optical pumping efficiency $\alpha \to 1$, and neglect the unwanted coupling of the control field to the $|2\rangle \leftrightarrow |3\rangle$ transition and set the anti-Stokes field to zero. We set $\Omega_{12} = \Omega$, and transform to the retarded time coordinate $\tau = t - z/c$ such that

$$\begin{aligned} \partial_z + \frac{1}{c}\partial_t &= \partial_z|_\tau \\ \partial_t &= \partial_\tau|_z. \end{aligned} \tag{2.50}$$

This reduces Equations 2.47 – 2.49 to:

$$\partial_\tau B = -\mathrm{i}\sqrt{\frac{d\gamma}{L}}\frac{\Omega^*}{\Gamma_s}S - \frac{|\Omega|^2}{\Gamma_s}B \tag{2.51}$$

$$\partial_z S = +\mathrm{i}\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_s}B - \frac{d\gamma}{L}\frac{1}{\Gamma_s}S. \tag{2.52}$$

These equations are linear in S and B, and therefore can be characterised by a Green's function mapping:

$$B_{\mathrm{stor}}(z) = \int_{-\infty}^{\infty} \mathrm{d}\tau\, K_1(z,\tau) S_{\mathrm{in}}(\tau) \tag{2.53}$$

$$S_{\mathrm{out}}(\tau) = \int_0^L \mathrm{d}z\, K_2(z,\tau) B_{\mathrm{stor}}(z), \tag{2.54}$$

where $K_{1(2)}$ characterises the storage (retrieval) process and encodes information about the read-in(out) control pulse.

At this point it is useful to define the number of excitations as:

$$N_{\mathrm{in(out)}} = \int_{-\infty}^{\infty} \mathrm{d}\tau \left\langle S^\dagger_{\mathrm{in(out)}}(\tau) S_{\mathrm{in(out)}}(\tau) \right\rangle \tag{2.55}$$

$$N_B = \int_0^L \mathrm{d}z \left\langle B^\dagger_{\mathrm{stor}}(z) B_{\mathrm{stor}}(z) \right\rangle \tag{2.56}$$

and the storage, retrieval, and total memory efficiencies as:

$$\eta_{\mathrm{stor}} = N_B/N_{\mathrm{in}} \tag{2.57}$$

$$\eta_{\mathrm{ret}} = N_{\mathrm{out}}/N_B \tag{2.58}$$

$$\eta_{\mathrm{mem}} = N_{\mathrm{out}}/N_{\mathrm{in}} = \eta_{\mathrm{stor}}\eta_{\mathrm{ret}}. \tag{2.59}$$

If the detuning $\Delta \gg \delta_c, \Omega, \gamma$ then an analytic expression for the Green's function can be found [192], and the storage interaction is determined by:

$$K_1(z,t) = \mathrm{e}^{\epsilon(\tau)W/\Gamma_s} J_0\left[2C_s\sqrt{\epsilon(t)(1-z)}\right] \tag{2.60}$$

where J_0 is the zeroth order Bessel function, $C_s = \sqrt{W\,d\gamma}/\Gamma_s$ is the Raman coupling constant, and we have defined a "normalised time" coordinate:

$$\epsilon(\tau) = \frac{1}{W}\int_{-\infty}^{\tau} \mathrm{d}\tau' \left|\Omega(\tau')\right|^2 \tag{2.61}$$

with normalisation

$$W = \int_{-\infty}^{\infty} \mathrm{d}\tau' \left|\Omega(\tau')\right|^2 \tag{2.62}$$

i.e. the integrated Rabi frequency. In general, an analytic expression for the Green's function cannot be found but it can be constructed numerically by solving the equations of motion.

We can perform a singular value decomposition (SVD) on $K_1(z,t)$ and express it as a convex sum of eigenmodes:

$$K_1(z,\tau) = \sum_k \psi_k(z)\,\lambda_k\,\phi_k^*(\tau) \tag{2.63}$$

where $\{\psi_k(z)\}$ and $\{\phi_k(\tau)\}$ are sets of orthonormal modes in the space of spin-waves and input temporal modes respectively [193]. The singular values λ_k map an input mode ϕ_k to a stored mode ψ_k with probability λ_k^2. The optimal storage efficiency is given by the square of the largest singular value $\eta_{\mathrm{stor,opt}} = \lambda_1^2$ and the optimal input mode is the corresponding $\phi_1(\tau)$. It can be shown that the optical depth is the only parameter that determines the *optimal* storage efficiency [192], which in the limit of large optical depths is given by $\eta_{\mathrm{stor,opt}} \approx 1 - 2.9/d$ [189]. For large optical depths the Raman memory can therefore, in principle at least, reach unit efficiency.

The number of modes that can be stored in the memory, N_{mode}, is given by the number of non-negligible singular values, and it can be shown that $N_{\mathrm{mode}} \sim \sqrt{d}/3$ [193]. Therefore the coupling strength, C_s scales much faster in the optical depth, and for reasonable coupling strengths, $C_s \sim 1$, the Raman memory is essentially single mode. Only a single input temporal mode will be stored in the memory, and all modes orthogonal to this will be transmitted straight through the memory virtually unaffected.

This modal capacity is a feature of the Raman memory that has thus-far been unexplored experimentally, and is investigated in Chapter 6.

2.2.3 Beam Splitter Memory Interaction

We can make a series of coordinate transformations to express the ideal Raman memory equations in a more tangible form.

We use dimensionless units: $\tilde{\tau} = \gamma\tau$, $\tilde{z} = z/L$, such that our variables are transformed to:

$$\widetilde{S} = S/\sqrt{\gamma}, \quad \widetilde{\Omega} = \Omega/\gamma, \quad \widetilde{\Gamma}_s = \Gamma_s/\gamma, \quad \widetilde{B} = B\sqrt{L}. \tag{2.64}$$

This transforms the ideal memory Equations 2.51 and 2.52 into:

$$\left(\partial_{\tilde{\tau}} + \frac{|\widetilde{\Omega}|^2}{\widetilde{\Gamma}_s}\right)\widetilde{B} = -\mathrm{i}\frac{\widetilde{\Omega}^*\sqrt{d}}{\widetilde{\Gamma}_s}\widetilde{S} \tag{2.65}$$

$$\left(\partial_{\tilde{z}} + \frac{d}{\widetilde{\Gamma}_s}\right)\widetilde{S} = +\mathrm{i}\frac{\widetilde{\Omega}\sqrt{d}}{\widetilde{\Gamma}_s}\widetilde{B}. \tag{2.66}$$

Again we introduce a normalised time coordinate

$$\epsilon(\tilde{\tau}) = \frac{1}{\widetilde{W}}\int_{-\infty}^{\tilde{\tau}} \mathrm{d}\tilde{\tau}' \left|\widetilde{\Omega}(\tilde{\tau}')\right|^2 \tag{2.67}$$

with normalisation

$$\widetilde{W} = \int_{-\infty}^{\infty} \mathrm{d}\tilde{\tau}' \left|\widetilde{\Omega}(\tilde{\tau}')\right|^2 = W/\gamma. \tag{2.68}$$

We can then perform a unitary transformation[3] on the signal and spin-wave:

$$\widetilde{S}(\tilde{z},\tilde{\tau}) = \frac{\widetilde{\Omega}(\tilde{\tau})}{\sqrt{\widetilde{W}}}\sigma(\tilde{z},\epsilon)\mathrm{e}^{-\left(\frac{\widetilde{W}\epsilon + d\tilde{z}}{\widetilde{\Gamma}_s}\right)} \tag{2.69}$$

$$\widetilde{B}(\tilde{z},\tilde{\tau}) = \beta(\tilde{z},\epsilon)\mathrm{e}^{-\left(\frac{\widetilde{W}\epsilon + d\tilde{z}}{\widetilde{\Gamma}_s}\right)}. \tag{2.70}$$

This transforms Equations 2.65 and 2.66 to:

$$\partial_{\tilde{z}}\sigma = -\mathrm{i}\widetilde{C}_s\beta \tag{2.71}$$

$$\partial_{\omega}\beta = +\mathrm{i}\widetilde{C}_s\sigma \tag{2.72}$$

where $\widetilde{C}_s = \sqrt{\widetilde{W}d}/\widetilde{\Gamma}_s$ or $C_s = \sqrt{Wd\gamma}/\Gamma_s$ is the Raman coupling constant for the signal mode. In the far detuned limit $\Gamma_s \approx \mathrm{i}\Delta$ and Equations 2.71 and 2.72 have the form of a beam splitter between modes σ and β. This demonstrates that, in the adiabatic limit, the ideal Raman memory equations are analogous to a light-matter beam splitter between the propagating signal mode and the stationary spin-wave mode, and the Hamiltonian for the system can be written as:

$$\mathcal{H}_s \propto C_s\hat{\sigma}\hat{\beta}^\dagger + \mathrm{h.c.} \tag{2.73}$$

The "reflectivity" of the beam splitter, or the coupling efficiency between modes σ and β is determined by the Raman coupling constant C_s. The storage and retrieval interactions can both be described as a beam splitter interaction with efficiency C_s^2, and therefore the total memory efficiency $\eta_{\mathrm{mem}} \propto C_s^4$.

It can be shown that the optimal memory efficiency increases quickly for small values of C_s but then plateaus for $C_s \sim 2$, and in order to have an efficient memory interaction we require $C_s^2 \gtrsim 1$ [167, 189]. If we assume the control pulse is a top-hat function with maximum Rabi

[3] This transformation is unitary in the adiabatic limit of large detunings, $\Delta \gg \Omega$.

frequency $\Omega_{\max}$ and bandwidth δ_c then $W = \int_{-\infty}^{\infty} \mathrm{d}\tau \ |\Omega(\tau)|^2 = |\Omega_{\max}|^2 / \delta_c$ and

$$|C_s|^2 \approx \left(\frac{d\gamma}{\delta_c}\right) \left|\frac{\Omega_{\max}}{\Delta}\right|^2 \tag{2.74}$$

and therefore in the adiabatic approximation where $\Omega/\Delta \ll 1$ we require $(d\gamma/\delta_c) \gg 1$ for an efficient memory protocol. In typical operational conditions we have $\delta_c \sim 1$ GHz and $\gamma \sim 5$ MHz, and therefore we require an optical depth $d \gg 200$ to achieve a high memory efficiency.

Equation 2.74 shows that the efficiency of the Raman memory protocol is proportional to $d^2(\Omega/\Delta)^4$, which we will discuss further in Section 2.3.1.

2.2.4 Four-wave mixing Noise

To explore the noise interaction we again consider Equations 2.47 – 2.49 but this time with no input signal field, $S = 0$:

$$\partial_t B = \mathrm{i}\sqrt{\frac{d\gamma}{L}}\frac{\Omega^*}{\Gamma_a} A^\dagger - \frac{|\Omega|^2}{\Gamma_a^*} B \tag{2.75}$$

$$\left(\partial_z + \frac{1}{c}\partial_t\right) A = \mathrm{i}\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_a} B^\dagger - \frac{d\gamma}{L}\frac{1}{\Gamma_a^+} A \tag{2.76}$$

where $\Gamma_a^+ = \Gamma_a + \mathrm{i}\Delta_{\mathrm{hf}}$. In normalised coordinates, with $\tilde{A} = A/\sqrt{\gamma}$, $\widetilde{\Gamma}_a^+ = \Gamma_a^+/\gamma$ this becomes:

$$\left(\partial_{\tilde{\tau}} + \frac{|\widetilde{\Omega}|^2}{\widetilde{\Gamma}_a^*}\right) \widetilde{B} = +\mathrm{i}\frac{\widetilde{\Omega}^*\sqrt{d}}{\widetilde{\Gamma}_a} \widetilde{A}^\dagger \tag{2.77}$$

$$\left(\partial_{\tilde{z}} + \frac{d}{(\widetilde{\Gamma}_a^+)^*}\right) \widetilde{A}^\dagger = -\mathrm{i}\frac{\widetilde{\Omega}^*\sqrt{d}}{\widetilde{\Gamma}_a^*} \widetilde{B}. \tag{2.78}$$

We perform a unitary transformation on the anti-Stokes and spin-wave:

$$\widetilde{A}^\dagger(\tilde{z}, \tilde{\tau}) = \frac{\widetilde{\Omega}^*(\tilde{\tau})}{\sqrt{\widetilde{W}}} \alpha^\dagger(\tilde{z}, \epsilon) \mathrm{e}^{-\left(\frac{\widetilde{W}\epsilon}{\widetilde{\Gamma}_a^*} + \frac{d\tilde{z}}{(\widetilde{\Gamma}_a^+)^*}\right)} \tag{2.79}$$

$$\widetilde{B}(\tilde{z}, \tilde{\tau}) = \beta(\tilde{z}, \epsilon) \mathrm{e}^{-\left(\frac{\widetilde{W}\epsilon}{\widetilde{\Gamma}_a^*} + \frac{d\tilde{z}}{(\widetilde{\Gamma}_a^+)^*}\right)} \tag{2.80}$$

which yields:

$$\partial_{\tilde{z}} \alpha^\dagger = -\mathrm{i}\widetilde{C}_a^* \beta \tag{2.81}$$

$$\partial_\omega \beta = +\mathrm{i}\widetilde{C}_a \alpha^\dagger \tag{2.82}$$

where $C_a = \sqrt{W d\gamma}/\Gamma_a = C_s \times (\Gamma_s/\Gamma_a)$. We can see that the anti-Stokes interaction can be described as a two-mode squeezing Hamiltonian between the anti-Stokes mode, α and the spin-wave, β as:

$$\mathcal{H}_a \propto C_a \hat{\alpha}^\dagger \hat{\beta}^\dagger + \text{h.c.}. \tag{2.83}$$

The four-wave mixing noise process consists of spontaneous anti-Stokes scattering, described by a two-mode squeezing Hamiltonian $\mathcal{H}_a$, followed by the desired memory retrieval interaction described by the beam-splitter Hamiltonian $\mathcal{H}_s$. The total efficiency of the process is proportional to $C_s^2 C_a^2$, and hence to maximise the signal to noise ratio we want to enhance C_s and suppress C_a.

This theoretical framework allows us to describe the Raman memory protocol and the four-wave mixing noise process, and determine how the efficiency and noise scale with experimental parameters.

2.3 Thesis Outline

In this thesis I investigate methods to increase the signal to noise ratio of the Raman memory. We consider this in two halves: firstly how to maximise the memory efficiency, and secondly how to suppress unwanted noise processes. Methods to achieve both of these goals are outlined below, and explored in the following chapters.

2.3.1 Increasing the Memory Efficiency

We saw above that the memory efficiency, η_{mem}, is proportional to C_s^4, and according to Equation 2.74 the Raman coupling constant, C_s, scales approximately as:

$$C_s \propto \sqrt{d}\left(\frac{\Omega}{\Delta}\right). \tag{2.84}$$

To increase the memory efficiency we want to maximise C_s, and we therefore need to either increase the optical depth, d, or the control pulse energy, Ω, or decrease the detuning, Δ, and these are explored more below.

1. **Decrease the detuning, Δ.**
   Decreasing the detuning from resonance would increase the Raman coupling strength, but also leads to more linear absorption of the signal and control fields. Absorption of the signal reduces the overall efficiency since this signal cannot be retrieved, and absorption of the control field reduces the pulse energy which diminishes the memory efficiency, and also leads to fluorescence from the excited state which can pollute the retrieved field.

   The caesium absorption profile is broadened by collisions with a buffer gas, which is added to the atomic vapour to inhibit motion out of the interaction region and increase the memory storage time. The memory currently operates at $\Delta = 15.2$ GHz which offers a good compromise of high coupling strength and low linear absorption, given the current operational temperature and buffer gas pressure. Reducing the pressure of the buffer gas to decrease the collisional broadening would enable a smaller operational detuning, but it would also decrease the time for caesium atoms to diffuse out of the interaction region and therefore affect the memory lifetime. There is an inherent compromise between spatial

confinement of the caesium atoms via buffer gas collisions to enable long storage times, and minimising the pressure broadening of the excited state to decrease the detuning and maximise the memory efficiency.

Alternatively, one could remove the buffer gas entirely and increase the optical beam size for a large interaction region. To achieve a long storage time a vapour cell with spin-preserving coating is also required such that collisions of the atoms with the walls of the cell do not change their internal state. This enables long storage times to be reached without collisional broadening due to a buffer gas, so the operational detuning of the memory can be reduced. However, as we shall see in Chapter 3 the buffer gas plays an important role in optical pumping to initialise the ensemble in a single ground state. Without a buffer gas it is challenging to reach high optical depths of the ensemble, which limits the achievable efficiency.

2. **Increase the optical depth, d.**
   The optical depth, d, of an atomic ensemble increases exponentially with the temperature of the atomic vapour, and therefore we wish to operate the memory as warm as possible. However, as we will discuss in Chapter 3, the fidelity with which we can initialise the entire ensemble in a single ground state degrades with temperature and sets an upper limit on the operational temperature. In Chapter 3 I investigate methods to achieve a high degree of spin polarisation in a warm atomic vapour at higher temperatures, to increase the Raman coupling strength.

3. **Increase the Rabi frequency of the control field, Ω.**
   We require a strong Rabi frequency to drive the memory interaction, and hence use a high energy pulsed laser. However, as the Rabi frequency increases, it induces an AC Stark shift between states $|2\rangle$ and $|3\rangle$, perturbing the two-photon resonance condition. This causes the memory efficiency to decrease above a threshold control Rabi frequency when the AC Stark shift, $(\Delta E)_{\mathrm{ACS}} \propto \Omega^2/\Delta$, is approximately equal to the bandwidth of the pulses, $\delta_{c,s}$. It is known that we can compensate for this by adding a dynamically varying phase to the control pulse to counteract the AC Stark shift [180]. In Chapter 6 I explore how shaping the control field can maximise the light-matter coupling strength and enable higher memory efficiencies.

2.3.2 Suppressing Four-Wave Mixing Noise

We discussed above in Section 2.1 that the Raman memory suffers from four-wave mixing noise which pollutes the photon number statistics of the retrieved state. It is critical to suppress this noise in order to have high fidelity storage and retrieval of photonic states.

There are several potential methods we can consider to suppress the unwanted two-mode squeezing interaction:

1. **Selection Rules**
   It is possible to prohibit four-wave mixing using atomic selection rules for a near-resonant memory in warm vapour [194] or solid state systems [195]. A magnetic field is applied

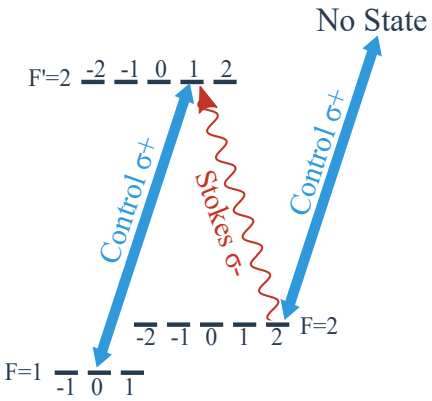


Figure 2.3: An example of how selection rules can be used to suppress four-wave mixing in a lambda energy level system, using the specific example of the D_1 line in rubidium. Here the initial and storage states are $|F = 2\rangle$ and $|F = 1\rangle$ ground states respectively, and the excited state has $|F' = 2\rangle$. If the atoms are initialised in $|F = 2, m_F = 2\rangle$ then the signal with σ_- polarisation couples to $|F' = 2, m'_F = 1\rangle$, and the control field with σ_+ polarisation completes the Raman transition to the $|F = 1, m_F = 0\rangle$ ground state. The control field cannot couple to the initial state since there is no excited state with $m'_F = 3$.

such that the Zeeman sublevels of the ground state, $|F, m_F\rangle$, are no longer degenerate and the population is optically pumped into the largest m_F state. The signal and control fields have circular polarisation σ_- and σ_+ respectively, such that the memory interaction decreases m_F by two and couples the initial state $|F, m_F\rangle$ to the storage state $|F', m'_F = m_F - 2\rangle$, as shown in Figure 2.3. The control field with σ_+ polarisation cannot couple to the initial state as no excited state with $m'_F = m_F + 1$ exists, and hence the unwanted coupling is prohibited. However, it has been shown that this method does not work for off-resonant memories in alkali vapours due to the destructive interference of different pathways for the memory interaction [196].

2. **Density of states**

   It is possible to suppress four-wave mixing noise by reducing the density of states that the anti-Stokes photon can be emitted into. In Chapter 4 I investigate operating the Raman memory inside a low-finesse optical cavity which reduces the density of scattering states at the anti-Stokes frequency, thereby inhibiting unwanted anti-Stokes scattering. Since the cavity is on-resonance for the control field, this has the added benefit of enhancing the Raman coupling constant and increasing the memory efficiency for a given control energy. However, it is technically challenging to build a cavity with the requisite resonance criteria with low intra-cavity loss.

3. **Phase Matching**

   The four-wave mixing process involves the absorption of two control photons to generate a signal and an anti-Stokes photon, and this process must satisfy both energy and momentum conservation (phase matching):

$$2\omega_c - \omega_s - \omega_a = 0, \tag{2.85}$$

$$2\boldsymbol{k}_c + \boldsymbol{k}_s + \boldsymbol{k}_a = 0. \tag{2.86}$$

The signal and anti-Stokes fields experience different dispersion in the atomic medium, and therefore it is not possible to satisfy the phase matching condition if the beams are all co-linear. The phase mismatch is given by

$$\begin{aligned} \delta k &= 2k_c - k_s - k_a && (2.87) \\ &\approx \frac{d\gamma}{L}\left(\frac{2}{\Delta+\Delta_{\mathrm{hf}}} - \frac{1}{\Delta} - \frac{1}{\Delta+2\Delta_{\mathrm{hf}}}\right) && (2.88) \end{aligned}$$

and hence to increase δk in order to suppress four-wave mixing we want to minimise Δ. However, as we discussed above, operating closer to resonance leads to linear absorption of the signal field and increases fluorescence noise. In Chapter 5 I investigate suppressing four-wave mixing by introducing an absorption feature at the frequency of the noise field to destroy the phase matching condition.

2.3.3 Exploring Temporal Mode Properties

A second goal of this thesis is to explore the temporal mode properties of the Raman memory. We discussed in Section 2.2.2 that, in the low coupling regime, the Raman memory is essentially single mode. In Chapter 6 I investigate the modal capacity of the Raman memory, and quantify the degree to which the memory is single mode. I also explore how the Raman memory can be used as a device to separate, delay, and manipulate temporal modes of light. The infrastructure developed in this chapter will also enable shaping of the control field to counteract the AC Stark shift and increase the memory efficiency, although this has not been demonstrated yet.

2.4 Summary

In this chapter I have introduced the theoretical framework of the Raman memory. I have re-derived the equations of motion that govern the system, including the desired beam-splitter Hamiltonian and the unwanted two-mode squeezing Hamiltonian, and described how the ideal Raman memory is single-mode. The rest of the thesis will explore the ideas of increasing the signal to noise ratio of the Raman memory, and investigating the temporal mode properties.

- In Chapter 3 I investigate methods to increase the memory efficiency by increasing the optical depth of the ensemble.
- In Chapter 4 I investigate suppressing four-wave mixing noise by operating the memory in an optical cavity which suppresses the density of states at the anti-Stokes frequency.
- In Chapter 5 I investigate suppressing four-wave mixing noise by introducing an absorption feature at the frequency of the noise field, thereby increasing the phase mismatch.
- In Chapter 6 I investigate the temporal mode properties of the Raman memory, and how the memory can separate and manipulate temporal modes of light.

Chapter 3:

High Gain Raman Memory

In this chapter I investigate methods to increase the efficiency of the Raman memory. We discussed in Chapter 2 that to increase the Raman coupling constant we need to either increase the optical depth, d, or the control Rabi frequency, Ω, or decrease the detuning, Δ. However, increasing the control pulse energy induces an AC Stark shift which perturbs the two-photon resonance condition of the memory, and decreasing the detuning introduces more linear absorption of the signal and control fields. We therefore investigate methods to increase the operational optical depth of the atomic ensemble.

We shall see below that the optical depth d of an atomic ensemble increases exponentially with temperature and therefore we want to operate the memory at as high a temperature as possible. However, as we will discuss in Section 3.2, the ability to initialise the ensemble in a single ground state degrades as we increase the temperature which causes a steep decline in memory performance above 70°C. In this chapter I will first describe the experimental infrastructure for implementing the Raman memory, and then present results on improving the spin-polarisation of caesium vapour at high temperatures to increase the operational optical depth.

The theory of optical pumping, and the results outlined in section 3.2 were presented in my Master's thesis [197] but are briefly repeated here for clarity. The results on the high gain Raman memory in Section 3.3 were performed during my PhD and are published in [198]. This project was overseen by D.J.S, and the experiment was performed by myself with assistance from D.J.S, P.M.L. and J.D.H.M. The numerical simulations were performed by J.N.

3.1 Experimental Set-up of the Raman Memory

There are three main elements for the experimental demonstration of a Raman memory: the preparation of the signal and control fields, the storage medium, and the filtering and detection of the retrieved signal field. A schematic for the experimental set-up for the storage of coherent states in the Raman memory is shown in Figure 3.2, and more details about these three stages

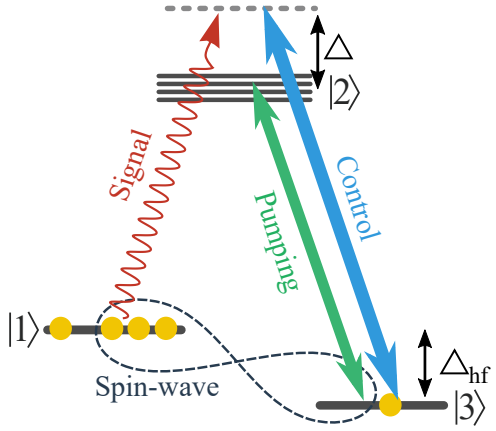


Figure 3.1: Schematic of the Raman memory.

are given below. The Raman memory was first demonstrated in 2010 [177], and has undergone several upgrades and refinements over the last decade. Whilst the key building blocks remain the same, I have fully rebuilt the set-up to improve certain elements. In this section I will outline the experimental architecture of the Raman memory and describe the adaptations I have made.

3.1.1 Storage Medium

We implement the Raman memory using the Λ system of the D_2 line $\left(6^2S_{1/2} \rightarrow 6^2P_{3/2}\right)$ in warm caesium $\left(^{133}Cs\right)$ vapour, and a schematic of the protocol is shown in Figure 3.1. The hyperfine ground states $6^2S_{1/2}$ $F = 4$ and $F = 3$ act as the initial state, $|1\rangle$, and the storage state, $|3\rangle$, respectively, and transitions between these long-lived states are dipole forbidden, enabling long coherence times. The entire hyperfine manifold of $6^2P_{3/2}$ with $F' = 2, 3, 4$ and 5 acts as the excited state $|2\rangle$, since these hyperfine states are separated by approximately 200 MHz and are not resolved due to Doppler broadening.

The motivation for choosing caesium is twofold: firstly it has the largest ground state splitting of the stable alkali atoms, $\Delta_{hf} = 9.2$ GHz. This separation sets an upper limit on the control pulse bandwidth such that the strong field addresses each ground state individually, and a large Δ_{hf} means that broadband signals can be stored. Secondly, caesium is the heaviest stable alkali which means that its diffusion time is the longest, and this slow diffusion out of the interaction region enables longer storage times. To further increase the lifetime of the memory we add a buffer gas to the atomic ensemble, to inhibit motion of the caesium atoms via spin-preserving collisions. As we see in Section 3.2 the choice of buffer gas plays an important role, and we discuss this in detail below.

The caesium vapour and buffer gas are contained in a cylindrical glass cell which is heated to approximately 70°C. The optical depth, d, is directly proportional to the number density of the atomic ensemble, $n(T)$, which is related to the vapour pressure, $P_v(T)$, via the ideal gas law: $n(T) = P_v(T)V/RT$. An empirical expression for the vapour pressure is given in [199], and it increases exponentially with temperature. Heating the vapour cell therefore increases the optical depth of the ensemble and the Raman light-matter coupling constant, although there

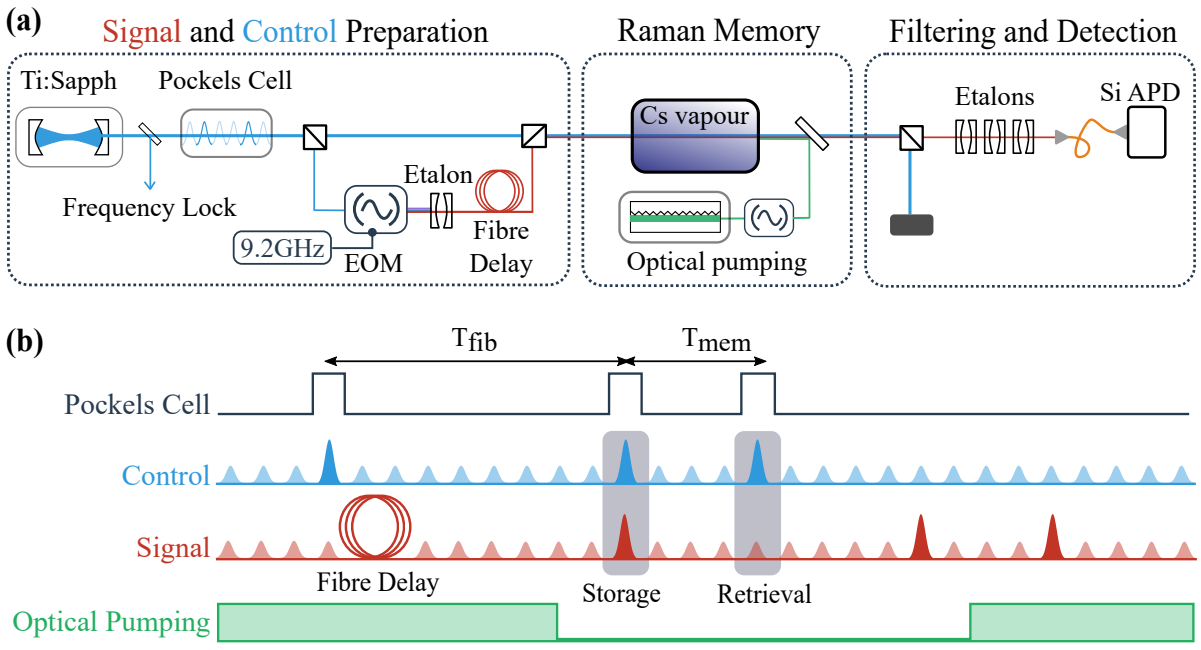


Figure 3.2: (a) A schematic of the experimental set-up of the Raman memory. The signal and control fields are derived from a mode-locked pulsed Titanium Sapphire laser (Ti:Sapph) which is frequency locked with respect to the caesium D_2 atomic line. Three pulses are picked from the train of pulses using a Pockels cell, and these act as the signal, the read-in and read-out control pulses. A small portion of the light passes through an electro-optic modulator (EOM) modulated at 9.2 GHz to generate a sideband at the signal frequency, and this frequency is isolated using a Fabry-Pérot etalon. The signal is coupled into a long single-mode fibre before being recombined with the control field on a polarising beam splitter (PBS) and focused in the centre of a caesium vapour cell. The atomic ensemble is initialised in $|1\rangle$ via a counter-propagating optical pumping beam. The orthogonally polarised strong control field is filtered away from the retrieved signal using a calcite beam displacer and a series of etalon frequency filters. The signal is coupled into a single-mode fibre and detected on a Silicon avalanche-photodiode (Si APD). (b) A schematic of the pulse timing. The Pockels cell is used to rotate the polarisation of three pulses by applying a high voltage in three time windows. The first pulse acts as the input pulse, and is temporally overlapped with the second pulse in the control mode by choosing the time delay between the first two pulses, T_{fib} to match the fibre delay in the signal mode. The third pulse acts as the retrieval control pulse, and the time between the second and third pulses can be adjusted to change the storage time of the memory, T_{mem}. The first control pulse reaches the memory T_{fib} before the memory interaction, but the optical pumping is still switched on during this time. The optical pumping is turned off using a fibre based EOM just before the signal is stored in the memory, and switched back on after the signal has been retrieved.

are limitations to the operational temperature of the Raman memory as we will discuss below. The cell is heated using quad-twist wire to minimise the generated magnetic field. The vapour cell is placed inside a triple-layered μ-metal shield to reduce the residual magnetic field by a factor of $\sim 10^3$ and limit magnetic dephasing of the ground-state coherence.

The atomic ensemble is initialised in state $|1\rangle$ via optical pumping which is described in detail in Section 3.2. The optical pumping beam is on resonance with the D_1 transition of caesium $\left(6^2S_{1/2} \rightarrow 6^2P_{1/2}\right)$, and is derived from a distributed feedback diode laser (Toptica DFB). The linewidth of this laser is specified to be around 4 MHz which is significantly more narrowband than the Doppler- and pressure-broadened caesium ensemble ($\sim$ 500 MHz) and therefore only a small fraction of the caesium atoms are on resonance with the optical pumping laser at any one time. However, the collisions of the caesium atoms with the buffer gas re-distribute the velocity classes within the ensemble, and after a few hundred nanoseconds[1] all the atoms will be frequency shifted onto resonance with the DFB laser and optically pumped into the initial state $|1\rangle$.

The central frequency of the DFB laser is passively stable to within 10 MHz over several days and therefore does not need to be actively frequency locked. The optical pumping beam is counterpropagating with the signal and control fields to minimise the amount of pumping light that reaches the detectors after the memory. The optical pumping is switched off during the memory interaction, as the strong Rabi frequency depletes the spin-wave and reduces the coherence time to around 100 ns. We use a fibre-based electro-optic modulator (EOM, EOSpace) to switch off the optical pumping beam with a switching time of approximately 5ns and an optimal extinction ratio of 200:1 which is sufficient to enable storage times on the order of microseconds[2].

To ensure that the spin-polarisation of the ensemble remains high during the memory interaction we investigated how the spin-polarisation evolves once the optical pumping beam has been switched off, or the T_1 time of the optical pumping. We turned off the optical pumping beam using a fibre-based EOM, and probed the ensemble after a certain delay time, Δt, by switching the scanning probe beam on and off with an acousto-optic modulator (AOM). This experiment was performed by P.M.L. and D.J.S. and the data (shown in Figure 3.3) demonstrate that the spin-polarisation decays over a timescale of $\sim$ 50 μs. Therefore, provided that the storage time of the memory is significantly shorter than 50 μs, the spin-polarisation of the ensemble does not decay over the memory lifetime.

3.1.2 Signal and Control Field Preparation

To drive the Raman memory we require a high power control pulse that is in two-photon resonance with a weak signal pulse. We initially investigate the storage of weak coherent states

[1] The mean free path of a caesium atom in the buffer gas is typically around 6 μm, and the average relative velocity between the caesium and buffer gas atoms at 70°C is around 650 ms^{-1} leading to an average collision rate of approximately 10^8 Hz. The average change in velocity per collision is approximately 150 ms^{-1} (depending on the weight of the buffer gas). The caesium atoms will spectrally diffuse through the entire Doppler profile in a few tens of collisions or a few hundred nanoseconds, and therefore the entire ensemble can be optically pumped in less than 500 ns [197].

[2] This extinction ratio later degraded due to photorefractive damage of the EOM, limiting the storage time in later chapters to several hundred nanoseconds.

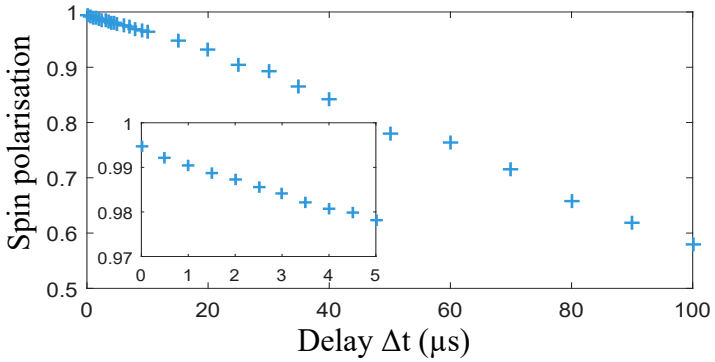


Figure 3.3: The spin-polarisation decay of the ensemble as a function of pump-probe delay Δt. The spin-polarisation is calculated by measuring the absorption spectrum of a scanning laser across the atomic transitions from which we can fit the absorption profile and infer the optical depth of the two transitions. More details of this method are given in [190]. The polarisation decays approximately linearly over a characteristic timescale on the order of 50 μs, and the inset shows the data for the first 5 μs. This experiment was performed by P.M.L. and D.J.S.

in the memory, but in Chapter 4 we discuss the storage of single photons.

The control and signal field pulses are both derived from the same mode-locked, pulsed Titanium-Sapphire (Ti:Sapph) laser. The Ti:Sapph is pumped by a continuous wave laser at 532 nm, which we upgraded to a Spectra-Physics Millennia eV to give a pump power of up to 15 W. The Ti:Sapph will remain stably mode-locked for a pump power of up to 12 W which gives a maximum Ti:Sapph output power of 2.5 W. The central wavelength of the Ti:Sapph is adjusted with a birefringent filter in the laser cavity, and is set to be 852 nm to be on resonance with the caesium D_2 transition. The bandwidth of the Ti:Sapph pulses is set by a Gires-Tournois interferometer (GTI) inside the laser cavity and is approximately 1.2 GHz, which is equivalent to a pulse duration of 360 ps. The frequency of the laser can be fine-tuned by changing the free-spectral range of the GTI, which allows adjustments over the mode-hop free range of approximately 40 GHz.

To frequency stabilise the laser with respect to the caesium atomic transition we use a frequency reference from an external-cavity diode laser (ECDL). The ECDL (Toptica DL Pro) is locked to a caesium hyperfine transition using Doppler-free saturated absorption spectroscopy. Here the beam is split into a weak forward-propagating probe beam and a strong counter-propagating pump beam, and the transmission of the weak beam through a room-temperature caesium vapour cell (with no buffer gas) is measured as it scans in frequency across the caesium D_2 atomic transitions. The transmission spectrum is shown in Figure 3.4, and two prominent absorption features are seen when the laser frequency is on resonance with the transitions from the two hyperfine ground states ($6^2S_{1/2}\,|F=3\rangle$ and $|F=4\rangle$) to the excited state ($6^2P_{3/2}$ hyperfine manifold). Within these absorption features we see structure which is due to the hyperfine levels of the excited state. When the laser is precisely on resonance with a transition from one of the ground states to one of the hyperfine levels in the excited state, for example $6^2S_{1/2}\,|F=3\rangle \rightarrow 6^2P_{3/2}\,|F'=2\rangle$, the strong pump beam saturates the absorption of the atoms and the weak probe will therefore stimulate emission from the excited state, which leads to an increase in transmission through the cell. We see this as a peak in the transmission through the

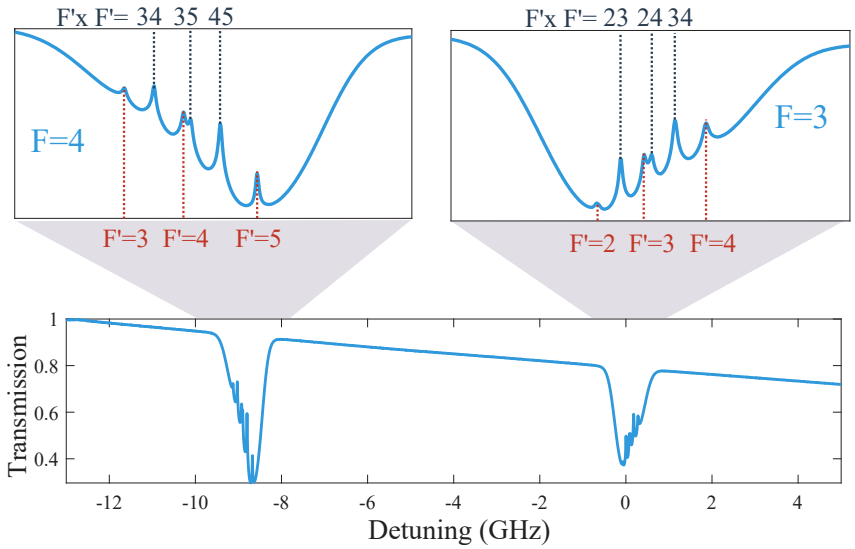


Figure 3.4: Frequency calibration using saturated absorption spectroscopy of a caesium vapour cell at room temperature. The transmission of a weak narrowband probe laser is measured through a caesium vapour cell as it scans in frequency across the atomic transitions. The absorption feature on the left (right) is the transition from the $6^2S_{1/2}\,|F=4\rangle$ ($|F=3\rangle$) ground state to the excited state manifold $6^2P_{3/2}\,|F'=2,3,4,5\rangle$. The overall gradient in transmission is due to the variation in output power of the narrowband laser as it scans in frequency. The upper plots show the zoomed-in frequency range around each transition, and we see peaks due to the hyperfine structure of the excited state, with $F'=2,3,4,5$, and the crossover peaks halfway between the hyperfine transitions, $F'\times F'$. These peaks are used as an exact frequency reference, and we define "zero" frequency offset to be the prominent crossover peak between the $6^2S_{1/2}\,|F=3\rangle \rightarrow 6^2P_{3/2}\,|F'=2\rangle$ and the $6^2S_{1/2}\,|F=3\rangle \rightarrow 6^2P_{3/2}\,|F'=3\rangle$ transitions.

cell when the frequency is exactly at an atomic resonance, visible above the Doppler-broadened absorption spectrum. These peaks are narrow and not subject to the Doppler broadening of the ensemble, as only the subset of atoms that have zero velocity along the direction of the beams will be on resonance with both beams simultaneously since these atoms experience no Doppler shift along the beam directions. There are also peaks in transmission halfway between the atomic transmissions known as crossover peaks. If the laser frequency is at the midpoint between the transitions to two excited state hyperfine levels, with frequencies f_1 and f_2, then the subset of atoms that experience a Doppler shift $\Delta f = \pm(f_1 - f_2)/2$ will be on resonance with the transition f_1 for one beam, and f_2 for the counterpropagating beam, and we see a transmission peak at these frequencies too. The exact frequency of these hyperfine and crossover transitions is well known [199], and we can use this as a precise frequency reference. The ECDL is frequency locked to the crossover peak between the $6^2S_{1/2}\,|F=3\rangle \rightarrow 6^2P_{3/2}\,|F'=2\rangle$ and the $6^2S_{1/2}\,|F=3\rangle \rightarrow 6^2P_{3/2}\,|F'=3\rangle$ transitions using this transmission spectrum via a Pound-Drever-Hall lock [200], and we use this as our "zero" frequency reference.

The locked ECDL and a small portion of the Ti:Sapph light are coupled into a scanning Fabry-Pérot interferometer (SFPI). The SPFI (Thorlabs SA210-8B) has a free-spectral range of 9.91 GHz, and can be used to compare the frequency of two fields within this range to a resolution of 67 MHz. The Ti:Sapph is brought onto resonance with the caesium D_2 transition by observing infra-red fluorescence through a caesium vapour cell, and then it is actively locked with respect to the ECDL reference by feeding back on the separation of the GTI plates inside the Ti:Sapph cavity. Once locked, the frequency can be manually shifted by tens of GHz from this set point, to allow detuning of the Raman memory with respect to the atomic transition. This frequency-offset lock was developed by D.J.S, and allows frequency stability of the Ti:Sapph on the order of 100 MHz over many hours.

The pulses from the Ti:Sapph have a repetition rate of 80 MHz, or a pulse separation of 12.5 ns. We select three pulses from this train to act as the signal, read-in and read-out control pulses using a fast Pockels cell. Applying a large voltage of to a Pockels cell rotates the polarisation of the light with a rise time of a few nanoseconds, and we can use it to rotate the polarisation of individual pulses by 90°. We pick three pulses from the pulse train with adjustable separations, and a calcite beam displacer after the Pockels cell directs the selected pulses towards the memory and discards the rest with a high extinction ratio. Previously, a Quantum Technology QC Pockels cell with a maximum repetition rate of 8 kHz was used, and this significantly limited the rate at which the experiment could be performed. We upgraded the Pockels cell and the high voltage switching electronics to a significantly faster set-up (Leysop RTP-X-4-20-AR650-1000-DMP with BME high voltage switching electronics) which can operate at a repetition rate of up to 2 MHz and has a switching time of 3 ns. The extinction ratio of this Pockels cell is also much higher and after fibre coupling the output mode we see an extinction ratio of 40,000:1.

The signal field needs to be red-detuned with respect to the control field by $\Delta_{\mathrm{hf}} = 9.2$ GHz to satisfy the two-photon resonance condition. To generate the signal mode, we split off a small portion of the Ti:Sapph light at a polarising beam splitter (PBS) and pass it through an electro-optic modulator (EOM) which is modulated at a frequency of 9.2 GHz. This generates

sidebands at ±9.2 GHz, and we use the red sideband as the signal frequency. We use a Fabry-Pérot etalon with a free-spectral range (FSR) of 36.8 GHz to filter away the carrier (control) frequency and the blue sideband, leaving only light at the signal frequency. This light is then coupled into a 80 m long single mode fibre (SMF), which delays the signal pulses by 400 ns (T_{fib}) with respect to the control field. A schematic of the pulse sequence is shown in Figure 3.2(b). The first of the three control pulses reaches the atoms T_{fib} before the memory interaction, but the optical pumping light is still switched on at this point so the ensemble is correctly initialised for storage. To ensure that the first pulse of the delayed signal mode (input signal) is temporally overlapped with the second pulse of the control arm (read-in), an adjustable delay is added to the control beam path to allow fine adjustments of the relative time delay. The time between the second and third pulses can be adjusted to change the storage time of the memory.

The selection rules of the Raman memory require that the signal and control fields are orthogonally polarised [189], and so the two modes are spatially recombined on a PBS. A half-waveplate on both the signal and the control beams before the PBS allows independent adjustment of the control pulse energy and the signal coherent state amplitude that is sent to the memory cell. The beams are collimated and then focused to a beam waist diameter of 280 μm in the centre of the caesium vapour cell.

3.1.3 Filtering and Detection

To drive the Raman memory we require a control pulse energy of around 500 pJ, which contains on the order of 10^{12} photons. To detect single-photon level signals without contamination from this strong control field it is therefore crucial that there is a high degree of filtering after the memory, and this is achieved using a combination of polarisation and spectral filtering. The signal and control fields have orthogonal polarisations and hence can be separated using a calcite beam displacer which has an extinction ratio of $\sim 10^5$. The signal is then coupled into a single-mode fibre before being passed through a series of air-spaced Fabry-Pérot etalons (FPEs) to further suppress the control field. The signal undergoes four passes of FPEs with an FSR of 18.4 GHz to maximally suppress the control field (which is 9.2 GHz separated from the signal), and also two passes through an etalon with FSR 103 GHz to suppress the anti-Stokes field (which is 18.4 GHz separated from the signal and therefore transmits through the first FPEs), in addition to further suppressing the control field.

I have upgraded the frequency filtering set-up to significantly improve the overall transmission and stability. Previously, the etalons were tuned onto resonance with the signal frequency by adjusting the angle of the etalon with respect to the incoming beam, which alters the effective cavity length and shifts the resonance frequency. However, the transmission through an etalon is maximum when the beam is at normal incidence to the planar etalon faces, as this maximises the interference of the field inside the cavity. We therefore now keep the etalons at normal incidence to the beam and tune the resonance frequency by adjusting the temperature. Each etalon is placed inside a mechanically stable mount, designed and built in-house, and we ensure the etalon is at normal incidence by measuring the transmission of the ECDL through the etalon as it scans in frequency. The transmitted signal is a series of Lorentzian peaks separated by the etalon FSR, and the position of the peaks moves as the angle of the

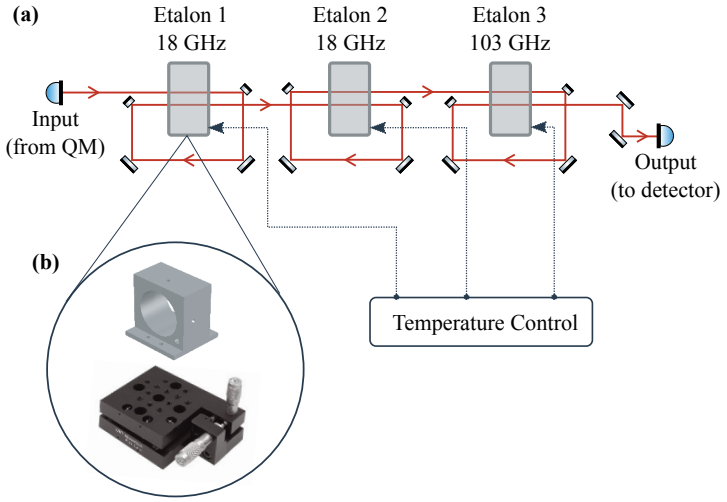


Figure 3.5: (a) Schematic of the experimental set-up of the frequency filtering using Fabry-Pérot etalons. The light from the quantum memory (QM) is transmitted through three etalons, two with a free-spectral range of 18 GHz and one with a free-spectral range of 103 GHz, in a double pass configuration. Small D-mirrors are used to individually address the two beams which are separated by approximately 5 mm. In between the two passes of each etalon, two mirrors are used to ensure that the two beams passing through the etalon are parallel. The output light is coupled into a single-mode fibre and sent to the detectors. (b) Each etalon is mounted in a mechanically stable, home-built mount which is temperature stabilised, on top of a two axis tilt stage (Newport M-36 Metric Tilt and Rotation Platform). By adjusting the two micrometers, the etalon can be brought to normal incidence with the incoming beam in both degrees of freedom.

etalon is adjusted. The transmission peaks shift down in frequency as the etalon becomes closer to normal and the effective path length increases. We adjust the angle of the etalon until the peaks are at the furthest red-detuned position, and then we temperature tune the etalons to shift the transmission peak to be at the signal frequency. The etalons are temperature tuned using precise temperature controllers (Oxford Instruments MercuryiTC) which provide active feedback to temperature stabilise the etalon mounts. A schematic of the filtering set-up is shown in Figure 3.5.

By temperature tuning the 18.4 GHz etalons, we can now achieve a transmission of the signal field of 90% for each pass compared to 65% for angle-tuning at room temperature. The transmission of the 103 GHz etalons is 70%, which is lower as we cannot heat the etalons enough to fully temperature tune them onto resonance, and have to use a combination of angle- and temperature-tuning. The retrieved signal from the memory is transmitted through four 18.4 GHz and two 103 GHz etalons, before being coupled into a SMF and sent to the detectors. The overall transmission of the filter system has been improved from around 0.5% to 10% due to temperature tuning, and the new filter design is also significantly more stable and stays aligned for several days compared to a few hours with the previous design. This factor of 20 improvement in the transmission of the filters leads to a 400-fold increase in detection rates for second-order autocorrelation measurements, and enables us to experimentally probe noise sources as we shall see later in Chapter 5.

Depending on the magnitude of the signal we detect it using either an avalanche photodiode

(APD) with a large maximum incident power (Menlo APD210), or a fibre-coupled single photon APD (Excelitas SPCM-AQRH). When using the Menlo we acquire data using a fast oscilloscope (Teledyne LeCroy wavesurfer 10), and when using the Excelitas APDs we use a timetagger (qutools quTAU) to generate an arrival time histogram. We measure a time trace in three configurations: firstly with just the signal field to measure the input signal; secondly with both the signal and control fields to implement the memory; and finally with just the control field and no signal to measure the noise generated by the control pulses. This set-up allows us to investigate storage of weak coherent states in the Raman memory, and measure the efficiency, lifetime and noise properties.

3.2 Optical Pumping

As we discussed at the beginning of this chapter, we wish to increase the optical depth of the atomic ensemble to increase the memory efficiency. We have seen that the optical depth d increases exponentially with the temperature of the atomic vapour, and therefore we want to operate the memory in warm caesium vapour. However, as we shall see below, the optical pumping efficiency degrades as we increase the temperature which causes a steep decline in memory performance above 70°C. In this section we therefore investigate how to achieve a high degree of spin-polarisation in a warm atomic vapour.

The first step of the Raman memory protocol is to prepare the atomic ensemble in a single ground state, $|1\rangle$. We define the spin-polarisation of the ensemble as $P = (N_1 - N_3)/(N_1 + N_3)$, where N_i is the number of atoms in state $|i\rangle$. An ensemble in thermal equilibrium with equal population in each ground state has $P = 0$, and an ideal spin-polarised ensemble with all atoms in $|1\rangle$ has $P = 1$. Spin polarisation is achieved through a process called optical pumping: a strong resonant beam is applied to the $|3\rangle \leftrightarrow |2\rangle$ transition (or indeed any transition from the ground state $|3\rangle$) to excite atoms out of the storage state $|3\rangle$. These atoms then decay from the excited state with approximately equal probability to the two ground states $|1\rangle$ and $|3\rangle$. After several cycles of excitation and decay, the atoms will accumulate in $|1\rangle$ and $P \to 1$. The excited state lifetime in caesium is 30.5 ns [199] and therefore the spin-polarisation reaches equilibrium on a time scale of approximately 200 ns.

However, the ability to optically pump an ensemble decreases as the temperature increases. When an atom decays from $|2\rangle \to |1\rangle$ it emits a photon that is on resonance with this transition which, if absorbed by another atom in the ensemble, can excite it out of the desired initial state $|1\rangle$ and "un-pump" the atom. As the optical depth of the atomic vapour increases, the probability of these resonant fluorescent photons escaping from the vapour without being absorbed by any other atoms rapidly approaches zero. The thermal background of resonant fluorescent photons can be absorbed and re-emitted multiple times, and "un-pump" many atoms in the vapour. This effect is known as radiation trapping [201], and significantly affects the ability to optically pump an atomic ensemble at high temperatures, limiting the achievable spin-polarisation. Previous demonstrations of the Raman memory used 20 Torr[3] of neon (Ne) buffer gas, in which the spin-polarisation of caesium degrades quickly after at temperatures

[3]Torr is a historic unit of pressure and 1 Torr is equivalent to 133.32 Pa or 1.33 mbar.

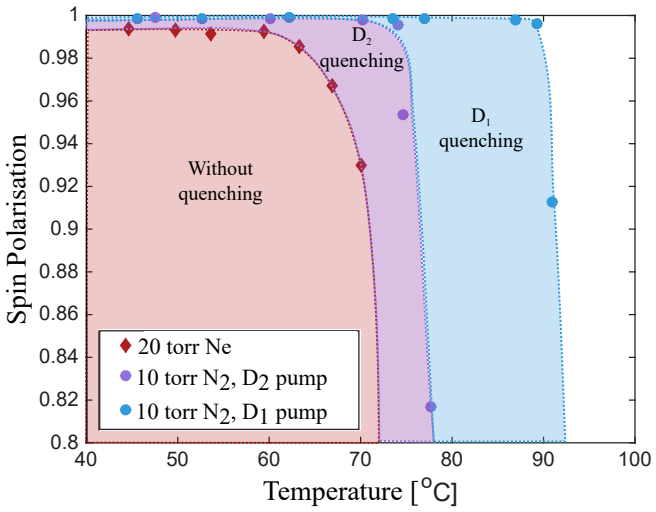


Figure 3.6: Measured spin-polarisation of warm caesium vapour in different buffer gases. The lines are to guide the eye and are not a fit to the data. Previous memory experiments used a buffer gas of 20 Torr of neon. Nitrogen has twice the collisional broadening and hence we compare 10 Torr of nitrogen to ensure the excited state broadening is the same, to operate the memory at the same detuning. We see that the spin-polarisation in nitrogen is significantly higher due to collisional quenching. We investigate pumping on the D_1 and the D_2 lines of caesium and see that the polarisation is much higher for the D_1 transition as the quenching cross section is higher. By using 10 Torr of N_2 and pumping on the D_1 line we can operate the memory at the same detuning from resonance but at a much higher temperature.

$T \gtrsim 70°C$ due to radiation trapping.

The problem of radiation trapping can be reduced by using a long, thin cell or a hollow-core fibre such that there is only high optical depth in one dimension, but it is challenging to reach high densities in such systems due to surface adsorption [202]. During my Master's project I investigated an alternative method to mitigate radiation trapping known as collisional quenching [197]. I found that by adding a molecular buffer gas such as nitrogen (N_2), we can reduce the detrimental effect of radiation trapping which had previously limited the temperature of the ensemble to $\sim 70°C$. Adding nitrogen suppresses radiation trapping through a process known as collisional quenching: caesium atoms in the excited state collide with the nitrogen and transfer their energy to the vibrational modes of the molecules; decaying to the ground state without resonant fluorescence [203]. This enables high spin-polarisation of the ensemble at much higher temperatures, as shown in Figure 3.6. Furthermore, we found that by optically pumping on the D_1 caesium transition rather than the D_2 transition, the quenching mechanism is much more effective and we can suppress radiation trapping further [198]. By adding 10 Torr of N_2 buffer gas we can now maintain high spin-polarisation of $P > 0.995$ at temperatures up to 90°C. The collisional broadening of caesium in a nitrogen buffer gas is 19.8 MHz/Torr compared to 9.81 MHz/Torr for neon [204], and therefore adding 10 Torr of N_2 introduces the same level of pressure broadening as 20 Torr of Ne and the amount of linear absorption of the signal is unchanged. Increasing the temperature from 70°C to 90°C equates to an increase in the optical depth, d, of a factor of 4. This allows us to operate the memory at much higher optical depths than before and explore a new regime with a very high Raman coupling constant.

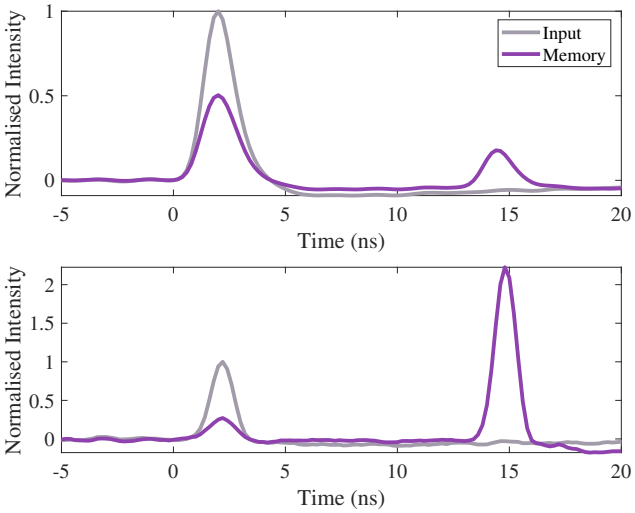


Figure 3.7: Typical arrival time traces for the Raman memory in the low gain (upper) and high gain (lower) regimes. The grey traces show the input signal, measured when the control field is turned off, and the purple traces show the signal when the control field is turned on and the signal is stored and then retrieved from the memory. The control Rabi frequency for both plots is $\Omega \sim 4.4$ GHz and the storage time is 12.5 ns. The temperature for the upper plot is $T = 72°\text{C}$ and the normalised output $\eta = 0.10$. For the lower plot the temperature is $T = 92°\text{C}$ and the four-wave mixing gain process amplifies the input signal such that the output is larger than the input, and $\eta = 1.81$.

As we discussed in Chapter 2, the Raman memory suffers from four-wave mixing (FWM) which introduces gain into the system. At high coupling strengths the gain interaction dominates and the Hamiltonian effectively becomes a two-mode squeezer which amplifies the input signal, and we can see artificially high memory efficiency. Whilst it is now clear that we need to suppress four-wave mixing by, for example, operating the Raman memory in a cavity to suppress this noise, the bulk (non-cavity) system is much easier to operate and so we can use it as a test bed for investigating the high coupling regime. The FWM gain process is well understood theoretically so we can use our results to indicate what the memory efficiency would be in the absence of four-wave mixing noise for the same Raman coupling strength.

3.3 High Temperature Raman Memory

Using the Raman memory set-up described above in Section 3.1, we investigate the storage of coherent states in this high Raman coupling regime enabled by improved optical pumping at high temperatures. Figure 3.7 shows the Raman memory traces measured using the Menlo APD and LeCroy oscilloscope in the low- and high-gain regimes, and we can see that the output pulse can be larger than the input. We define the normalised output, η, as the ratio of the pulse energy of the retrieved pulse and the input pulse. For the case of a noise-free memory, this quantity is simply the total efficiency of the memory protocol, η_{mem}. However, in the presence of FWM this quantity represents contributions from both parts of the Hamiltonian: the desired beam-splitter memory interaction, and FWM gain, and hence η can be larger than one.

We measure the normalised output, η, as a function of the control field Rabi frequency

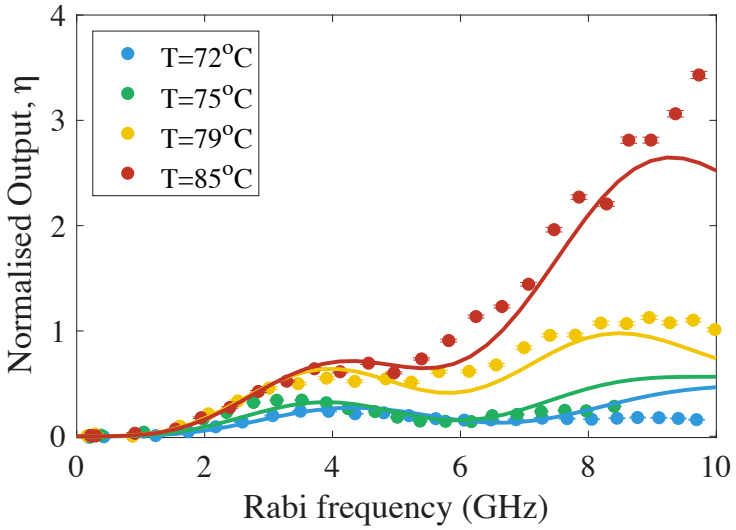


Figure 3.8: Normalised output, η, as a function of the control Rabi frequency for different temperatures of the caesium vapour, with detuning fixed at $\Delta = 15.2$ GHz. The solid lines are the predictions from numerical simulations, and agree well with the data for low Rabi frequencies before the adiabaticity condition breaks down at $\Omega \gtrsim 7$ GHz.

for different temperatures of the caesium ensemble, and the results are shown in Figure 3.8. For small control Rabi frequencies, $\Omega \lesssim 3$ GHz, η increases quickly with Ω since the Raman coupling strength C_s is proportional to $\sqrt{d}\Omega/\Delta$. However, as Ω increases, the strong control field induces an AC Stark shift between states $|2\rangle$ and $|3\rangle$, which perturbs the two-photon resonance condition and reduces the memory efficiency. At a detuning of $\Delta = 15.2$ GHz this occurs at $\Omega \approx 4$ GHz for each temperature, since the AC Stark shift, $(\Delta E)_{\mathrm{ACS}}$ which is proportional to Ω^2/Δ becomes comparable to the spectral bandwidth of the pulse, $\delta_{C,S} \sim 1.2$ GHz, and we see a decrease in η at this point. Nevertheless, at even higher Rabi frequencies we see that η increases dramatically, even far above one due to FWM gain. The FWM gain process increases approximately exponentially in Ω at high Rabi frequencies, and this gain interaction dominates and amplifies the input signal. We can see in Figure 3.8 that η is significantly greater at higher temperatures. This increase is a direct consequence of the higher Raman coupling strength due to increased optical depth which was enabled by improving the spin-polarisation at higher temperatures.

In addition to maximising the memory efficiency, we also need to consider the amount of linear absorption of the signal field, since if we operate in the single photon regime then it is crucial that there is a negligible amount of loss of the signal. As we increase the optical depth, the amount of linear absorption at a given detuning from resonance increases, as shown in Figure 3.9(a). When the operational temperature of the memory was 70°C, it was found that operating at $\Delta_0 = 15.2$ GHz was a good compromise for ensuring a high coupling strength but low linear absorption. As we increase the temperature it is clear that we need to increase the detuning if we wish to avoid additional linear absorption.

In Figure 3.9(b)–(c) we show the measured memory efficiency as a function of control field Rabi frequency for two different detunings: first at $\Delta_0 = 15.2$ GHz, and secondly at the detuning $\Delta_1(T)$ which gives the same absorption at temperature T as we get for Δ_0 at $T_0 = 70$°C. We can

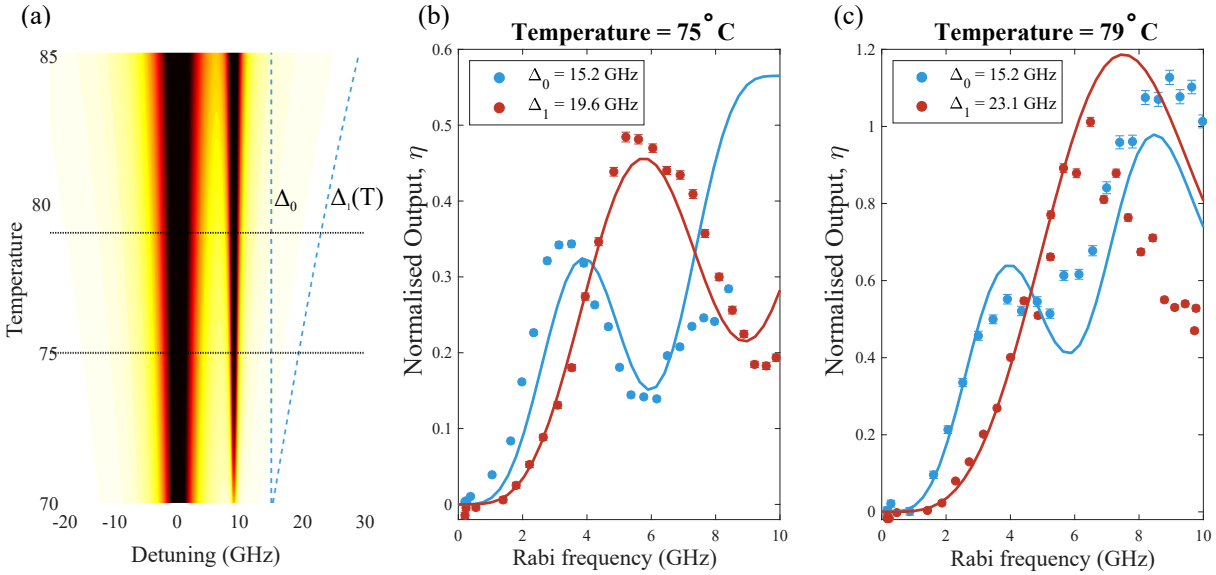


Figure 3.9: (a) Simulation of the linear absorption of a narrowband laser through caesium vapour as a function of detuning and temperature. The blue dashed lines show the detuning $\Delta_0 = 15.2$ GHz and $\Delta_1(T)$ which is the detuning at which the linear absorption at temperature T is the same as $\Delta_0(T = 70°\text{C})$ (b) The normalised output, η, as a function of the control field Rabi frequency for two different detunings: $\Delta_0 = 15.2$ GHz and $\Delta_1(T = 75°\text{C}) = 19.6$ GHz. (c) The normalised output as a function of Rabi frequency for $\Delta_0 = 15.2$ GHz and $\Delta_1(T = 79°\text{C}) = 23.1$ GHz.

see that at low Rabi frequencies, the efficiency is higher for Δ_0 as the Raman coupling strength is larger. When we operate at $\Delta_1(T)$, we decrease the coupling strength but also decrease the AC Stark shift, since this energy shift ΔE is proportional to Ω^2/Δ, and therefore the memory efficiency roll-off occurs at a higher Rabi frequency. This shows an additional benefit of the increased Raman interaction strength: the increase in optical depth means that we can operate further from resonance while still maintaining strong Raman coupling and therefore mitigate the AC Stark effect and achieve higher memory efficiency.

To model these data we numerically solve the linearised Maxwell-Bloch equations in one dimension, and these simulations were performed by J.N. The model assumes that the population of the excited state is negligible and that the population of the ground states remains constant throughout the interaction, since the average number of photons in the signal field is much smaller than the number of atoms in the ensemble. This model is valid in the adiabatic limit when the Rabi frequency, Ω, is significantly less than the detuning Δ. This theory has been developed by J.N. and has been used to model the Raman memory previously [205], but here we test it for the first time in a regime of high optical depth and high Rabi frequency.

Initial tests of the model predict a significantly higher memory efficiency than we measure and we believe this discrepancy is two-fold. Firstly, the residual optical pumping light left on during storage leads to an additional AC Stark shift that reduces the mode overlap between the signal and control fields. Secondly, due to the AC Stark shift from the control field, the central frequency of the retrieved signal pulse is shifted. This leads to a reduced transmission through our narrow etalon transmission function, and decreases the measured retrieval efficiency. We include these two features in our model by adding in a two-photon detuning term proportional

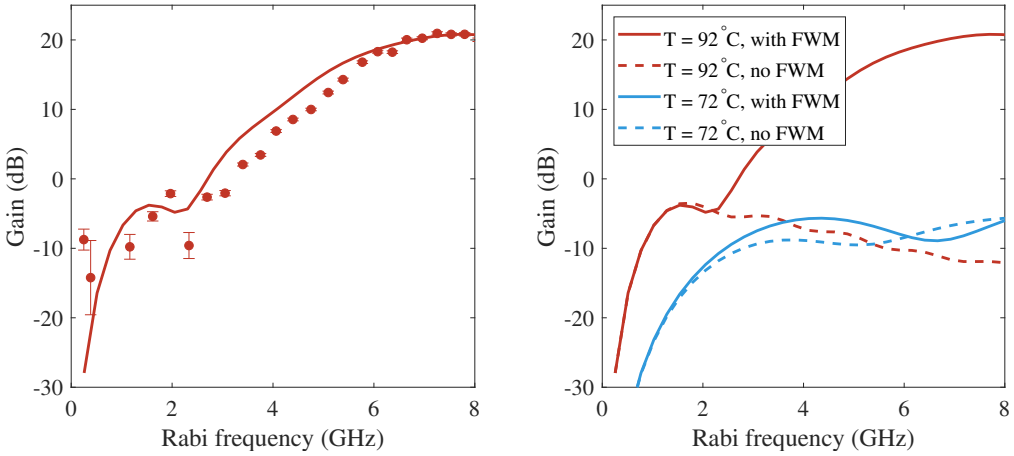


Figure 3.10: η in units of dB as a function of control field Rabi frequency for Cs vapour temperature of $\sim$ 92°C and detuning $\Delta = 15.2$ GHz. The dotted line shows the predicted memory efficiency when four-wave mixing is suppressed.

to the Rabi frequency of the pumping beam, and also by multiplying the retrieval efficiency by the transmission of the etalons we expect for that Stark shift at a given Rabi frequency.

With these additional features, we see a good agreement with our data. Figures 3.8 and 3.9 show that the model qualitatively captures the trends of the data for all temperatures, including the roll-over in efficiency due to the AC Stark shift, and the continued increase at higher Rabi frequencies due to four-wave mixing gain. We therefore have confidence that this linearised model can be extended to this high coupling regime. The discrepancies at very high Rabi frequencies reflect that the adiabatic approximation is no longer valid, and we see that this approximation breaks down when $\Omega \gtrsim \Delta/2$. In Figure 3.9(b) the model qualitatively agrees well with the data until $\Omega \sim 7$ GHz for a detuning of $\Delta = 15.2$ GHz, and until $\Omega \sim 9$ GHz for $\Delta = 19.6$ GHz, consistent with the conclusion of a breakdown of adiabaticity at very high Rabi frequencies. The model can therefore be used to predict performance in the Raman memory at high temperatures, provided that $\Omega \lesssim \Delta/2$.

The data in Figure 3.10 show the memory performance when the ensemble temperature is 92°C, and the normalised output, η, increases up to $\gtrsim 120$. This demonstrates extremely strong amplification of the input signal of over 20 dB due to the strong Raman interaction. These results could also be applied to the generation of strongly-squeezed light via four-wave mixing in alkali vapours [206].

We can use the model to predict what the memory performance would be in the absence of four-wave mixing noise. The dashed lines in Figure 3.10(b) shows the predicted memory efficiency if the four-wave mixing process was fully suppressed. The efficiency reaches a maximum of 45% for $T = 92$°C at $\Omega \sim 1.5$ GHz, compared to 27% for $T = 72$°C at $\Omega \sim 8$ GHz. Operating the memory at higher temperatures will allow for significantly higher efficiencies at much lower Rabi frequencies, even when four-wave mixing noise is suppressed.

3.4 Summary

In this chapter I have investigated increasing the Raman memory efficiency, and the key results are outlined below.

- In Section 3.1 I outlined the experimental architecture for the Raman memory, and discussed the upgrades I have made, in particular to the frequency filtering of the control field.
- In Section 3.2 I discussed that the main limiting factor for increasing the optical depth was the ability to optically pump the ensemble into a single ground state at high temperatures due to the detrimental effect of radiation trapping. I presented results on introducing a molecular buffer gas to the ensemble which reduces the effect of radiation trapping via collisional quenching, and enables high spin-polarisations of caesium at temperatures up to 90°C. This increases the optical depth by a factor of four compared to the previous operational temperature of 70°C, and enabled us to investigate the Raman memory in this high coupling strength regime.
- In Section 3.3 I investigated the high-coupling Raman memory by measuring the normalised output, η, as a function of the control Rabi frequency in different temperatures of the caesium vapour. I found that η was significantly higher at larger optical depths due to the stronger light-matter coupling. I also explored changing the detuning of the memory to counteract the AC Stark shift, and found that higher efficiencies could be reached at larger detunings, despite the lower Raman coupling constant. I presented the results of a numerical simulation of the Raman memory performed by J.N. that demonstrates strong agreement with our data. We used this model to predict what the memory performance would be in the absence of four-wave mixing noise, and saw that the increased optical depth will allow much larger efficiencies to be reached in a noise-free Raman memory.

Chapter 4:

Storage of Single Photons in a Cavity-Enhanced Raman Memory

In this chapter I will present results on the storage of single photons in a cavity-enhanced Raman memory, which is designed to suppress anti-Stokes scattering and inhibit four-wave mixing noise. The first demonstration of a cavity-enhanced Raman memory was conducted before my PhD and was led theoretically by J.N. and J.H.D.M, and experimentally by D.J.S, and these results are published in [205, 207]. I have since developed this work and built a next-generation cavity memory, with assistance from D.J.S. and P.M.L. I interfaced this with a single photon source which was built by B.B. and myself, and closely resembles the source presented in [183]. I will first present the theory of the cavity-enhanced Raman memory, which has been developed by J.N. and J.H.D.M, and then present current experimental results.

4.1 Theory of a Cavity-Enhanced Raman Memory

We discussed in Section 2.1 that the Raman memory suffers from four-wave mixing noise. In addition to the desired memory interaction the control field can also couple to the populated ground state and scatter an anti-Stokes photon, A, leading to the generation of a spurious spin-wave B. This spin-wave can be retrieved as a Stokes photon which cannot be filtered spectrally or temporally from the retrieved signal. This two-mode squeezing operation between A and B leads to spurious excitations in the signal mode, even when no input signal is present, and pollutes the photon number statistics of the retrieved state. Previously, single photons were stored in the Raman memory, but the second-order autocorrelation of the output was measured to be $g^{(2)}_{\text{out}} = 1.59 \pm 0.03$ [183], far above the ideal value of $g^{(2)} = 0$ due to four-wave mixing noise. It is critical to suppress this noise in order to have high fidelity storage and retrieval of photonic states.

In Section 2.3.2 we discussed different methods to suppress four-wave mixing noise which

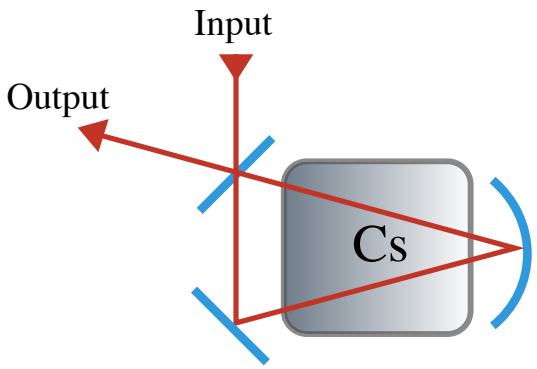


Figure 4.1: A ring cavity containing caesium atoms.

utilise selection rules, density of states and phase-matching. In this chapter we consider suppressing four-wave mixing noise by altering the *density of states* at the anti-Stokes frequency. We investigate operating the Raman memory inside a low-finesse optical cavity which is resonant for the signal and control fields and anti-resonant for the anti-Stokes field. This suppresses the density of scattering states at the anti-Stokes frequency, thereby inhibiting unwanted anti-Stokes scattering. We note that it is only necessary to suppress on-axis spontaneous anti-Stokes scattering (along the direction of propagation of the signal and control fields), as any spin-wave excitation produced by scattering in any other direction is not phase-matched for retrieval into the signal mode and therefore does not contribute to the noise. Therefore, a low-finesse cavity is sufficient for strong noise suppression.

The theory of noise suppression for a Raman memory in a cavity has been developed by J.N. and J.H.D.M, and is explored in detail in [190, 205]. An outline of this theory and the key results are summarised below.

The equations of motion for the Raman memory are given by Equations 2.47 – 2.49, which, in the case of perfect optical pumping ($\alpha \to 0$), reduce to:

$$(c\partial_z + \partial_t)\, S = \mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_s}B - \kappa_s S \tag{4.1}$$

$$(c\partial_z + \partial_t)\, A = \mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_a}B^\dagger - \kappa_a A \tag{4.2}$$

$$\partial_t B = -\mathrm{i}\sqrt{\frac{d\gamma}{L}}\left[\frac{\Omega^*}{\Gamma_s}S - \frac{\Omega}{\Gamma_a}A^\dagger\right] - \left[\frac{1}{\Gamma_a^*} + \frac{1}{\Gamma_s}\right]|\Omega|^2 B. \tag{4.3}$$

where

$$\kappa_s = \frac{d\gamma}{L}\frac{c}{\Gamma_s} \tag{4.4}$$

$$\kappa_a = \frac{d\gamma}{L}\frac{c}{\Gamma_a + \mathrm{i}\Delta_{\mathrm{hf}}}. \tag{4.5}$$

We implement the Raman memory inside a ring cavity, as shown in Figure 4.1. The cavity has a round-trip length of L, or equivalently a round-trip time $\tau = L/c$, and an input/output

coupler of reflectivity r (or transmissivity $t_r = \sqrt{1-r^2}$). The optical fields are subject to both dispersion and loss which can be quantified by

$$\zeta_{s,a} = \mu_{s,a} \mathrm{e}^{\mathrm{i}\phi_{s,a}} \tag{4.6}$$

where $\mu_{s,a} = r \exp\left(-\mathfrak{Re}\left\{\kappa_{s,a}\right\}\tau\right)$ is the round-trip amplitude transmission and accounts for loss from the in-coupler and the atomic ensemble, and $\phi_{s,a} = k_{s,a}L - \mathfrak{Im}\left\{\kappa_{s,a}\right\}\tau$ is the total phase accrued by the fields in one round trip including dispersion from the atomic ensemble.

We consider the intra-cavity field amplitudes, s, a and b, which are related to the incoming fields at the cavity in-coupler via:

$$s = \sqrt{\tau} S_0, \quad a = \sqrt{\tau} A_0, \quad b = \sqrt{L} B_0 \tag{4.7}$$

and define the cavity decay rates for each field as:

$$\frac{1}{\gamma_{s,a}} = \frac{\tau \zeta_{s,a}}{1-\zeta_{s,a}}. \tag{4.8}$$

The equations of motion are now given by:

$$\begin{aligned}
\partial_t s &= -\gamma_s s + \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_s} b + \frac{t_r}{\zeta_s\sqrt{\tau}} S_{\mathrm{in}} && (4.9)\\
\partial_t a &= -\gamma_a a + \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_a} b^\dagger + \frac{t_r}{\zeta_a\sqrt{\tau}} A_{\mathrm{in}} && (4.10)\\
\partial_t b &= -\mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\left[\frac{\Omega^*}{\Gamma_s} s - \frac{\Omega}{\Gamma_a} a^\dagger\right] - \left[\frac{1}{\Gamma_a^*} + \frac{1}{\Gamma_s}\right]|\Omega|^2 a. && (4.11)
\end{aligned}$$

We consider these equations in the so-called "bad-cavity limit", where the pulses are narrowband compared to the linewidth of the cavity, and $|\partial_t s| \ll |\gamma_s s|$. In this limit we can set $\partial_t s = \partial_t a \approx 0$, and solve to find:

$$\begin{aligned}
s &= \mathrm{i}c\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_s\gamma_s} b + \frac{t_r}{\gamma_s\zeta_s\sqrt{\tau}} S_{\mathrm{in}}, && (4.12)\\
a &= \mathrm{i}c\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_a\gamma_a} b^\dagger + \frac{t_r}{\gamma_a\zeta_a\sqrt{\tau}} A_{\mathrm{in}}, && (4.13)\\
\partial_t b &= \left\{\frac{d\gamma}{\tau}\left[\frac{1}{\Gamma_s^2\gamma_s} + \frac{1}{|\Gamma_a|^2\gamma_a^*}\right] - \left(\frac{1}{\Gamma_s} + \frac{1}{\Gamma_a^*}\right)\right\}|\Omega|^2 b + \mathrm{i}\frac{t_r\sqrt{d\gamma}}{\tau}\left[\frac{\Omega^*}{\zeta_s\Gamma_s\gamma_s} S_{\mathrm{in}} + \frac{\Omega}{\zeta_a^*\Gamma_a\gamma_a^*} A_{\mathrm{in}}^\dagger\right] && (4.14)
\end{aligned}$$

We make a coordinate transformation to remove the dependence on the temporal shape of the control field, by transforming to normalised time:

$$t \rightarrow \epsilon(t) = \frac{1}{W}\int_{-\infty}^{t} \mathrm{d}t' \ \left|\Omega(t')\right|^2 \tag{4.15}$$

$$s \rightarrow \sigma = s\frac{\sqrt{W}}{\Omega} \tag{4.16}$$

$$a \rightarrow \alpha = a\frac{\sqrt{W}}{\Omega} \tag{4.17}$$

where the normalisation constant is:

$$W = \int_{-\infty}^{\infty} \mathrm{d}t' \ \left|\Omega(t')\right|^2 . \tag{4.18}$$

Equations 4.12 - 4.14 then become:

$$\sigma = c_s b + p_s \sigma_{\mathrm{in}} \tag{4.19}$$

$$\alpha = c_a b^\dagger + p_a \alpha_{\mathrm{in}} \tag{4.20}$$

$$\partial_\epsilon b = fb + g_s \sigma_{\mathrm{in}} + g_a \alpha_{\mathrm{in}}^\dagger \tag{4.21}$$

where the coefficients are given by:

$$c_{s,a} = \mathrm{i}\sqrt{\frac{d\gamma W}{\tau}}\frac{1}{\Gamma_{s,a}\gamma_{s,a}} \tag{4.22}$$

$$p_{s,a} = = \frac{t_r}{\zeta_{s,a}\gamma_{s,a}\sqrt{\tau}} \tag{4.23}$$

$$f = W\left\{\frac{d\gamma}{\tau}\left[\frac{1}{\Gamma_s^2\gamma_s} + \frac{1}{|\Gamma_a|^2\gamma_a^*}\right] - \left(\frac{1}{\Gamma_s} + \frac{1}{\Gamma_a^*}\right)\right\} \tag{4.24}$$

$$g_s = \mathrm{i}\frac{t_r\sqrt{d\gamma W}}{\zeta_s\Gamma_s\gamma_s\tau}; \qquad g_a = \mathrm{i}\frac{t_r\sqrt{d\gamma W}}{\zeta_a^*\Gamma_a\gamma_a\tau}. \tag{4.25}$$

The parameter f quantifies the memory interaction and can be rewritten as

$$f = \left(C_s^2 + \frac{\gamma_s}{\gamma_a^*}\frac{\Gamma_a}{\Gamma_a^*}C_a^2\right) - W\left(\frac{1}{\Gamma_s} + \frac{1}{\Gamma_a^*}\right) \tag{4.26}$$

where

$$C_{s,a} = \frac{\sqrt{\mathcal{C}\gamma W}}{\Gamma_{s,a}} \tag{4.27}$$

and $\mathcal{C} = d\zeta_s/(1-\zeta_s)$ is the cooperativity of the cavity for the signal field. In the non-cavity case we had $C_{s,a} = \sqrt{d\gamma W}/\Gamma_{s,a}$ (Section 2.2.3) and therefore we identify $\mathcal{C}$ as the enhanced optical depth in the cavity. If the cavity is on-resonance for the signal field and $\phi_s = 0$, then the cooperativity, $\mathcal{C}$ is equivalent to the optical depth for a free-space medium of length $L \times \mathcal{F}_s/\pi$, where $\mathcal{F}_s$ is the finesse of the cavity for the signal field. Therefore there is an enhancement in the optical depth by the effective number of round trips of the cavity $\mathcal{F}_s/\pi$.

Solving Equations 4.19 – 4.21 gives:

$$\sigma_{\text{out}} = \frac{\mathrm{i}t_r\zeta_s\sqrt{d\gamma W}}{r\Gamma_s\gamma_s\tau}b + (\chi - r)\sigma_{\text{in}} \tag{4.28}$$

where the cavity transmission amplitude is defined as

$$\chi = \frac{t_r^2}{r}\frac{\zeta_s}{1-\zeta_s}. \tag{4.29}$$

If the system is initialised with no spin-wave (i.e. perfect optical pumping), then after the storage interaction we have [190, 205]:

$$\sigma_{\text{out},0}(\epsilon) = \int_0^1 \mathrm{d}\epsilon' \left[M_{\text{BS}}^{[0]}(\epsilon,\epsilon')\sigma_{\text{in},0}(\epsilon') + M_{\text{FWM}}^{[0]}(\epsilon,\epsilon')\alpha_{\text{in},0}^\dagger(\epsilon')\right]. \tag{4.30}$$

Here the superscript [0] indicates that the fields and integral kernels are those in the storage time bin. We can see that there are contributions from the beam-splitter (BS) and four-wave mixing (FWM) parts of the Hamiltonian, and these integral kernels are given by:

$$M_{\text{BS}}^{[0]}(\epsilon,\epsilon') = -\chi C_s^2 M_c(\epsilon,\epsilon') + (\chi - r)\delta(\epsilon - \epsilon') \tag{4.31}$$

$$M_{\text{FWM}}^{[0]}(\epsilon,\epsilon') = -\chi C_s C_a x M_c(\epsilon,\epsilon') \tag{4.32}$$

where

$$M_c(\epsilon,\epsilon') = \Theta(\epsilon,\epsilon')\exp\left[f(\epsilon-\epsilon')\right] \tag{4.33}$$

$$x = \frac{1-\zeta_s}{1-\zeta_a^*}. \tag{4.34}$$

Similarly the retrieved field is given by

$$\sigma_{\text{out},1}(\epsilon) = \int_0^1 \mathrm{d}\epsilon' \left\{M_{\text{BS}}^{[1]}(\epsilon,\epsilon')\sigma_{\text{in},0}(\epsilon') + M_{\text{FWM}}^{[1]}(\epsilon,\epsilon')\alpha_{\text{in},0}^\dagger(\epsilon') + M_{\text{FWM}}^{[0]}(\epsilon,\epsilon')\alpha_{\text{in},1}^\dagger(\epsilon')\right\} \tag{4.35}$$

where

$$M_{\text{BS}}^{[1]} = -\chi C_s^2 \mathrm{e}^f \mathrm{e}^{f(\epsilon-\epsilon')} \tag{4.36}$$

$$M_{\text{FWM}}^{[1]} = -\chi C_s C_a x \mathrm{e}^f \mathrm{e}^{f(\epsilon-\epsilon')}. \tag{4.37}$$

We notice that the coupling of the anti-Stokes field, C_a, always appears multiplied by x, and we identify it as the *noise-suppression factor*.

Using this framework, several key results can be derived [190, 205]:

1. **Memory Enhancement**

   The memory efficiency is given by

$$\eta_{\text{mem}} = \left|\chi C_s^2 E \kappa\right|^2 \tag{4.38}$$

   where $E = \int_0^1 \mathrm{d}\epsilon\, \mathrm{e}^{-\zeta\epsilon}$, $\zeta = -(f + f^*)$, and $\kappa = (\mathrm{e}^{\zeta} E)^{-1/2}\psi_{\text{in}}(1)$ quantifies the overlap between the input field, $\psi_{\text{in}}(\epsilon)$, and the cavity modes. In the coupling constant, C_s, the optical depth of ensemble d has been replaced by the cooperativity, $\mathcal{C} = d\mathcal{F}_s/\pi$ where $\mathcal{F}_s$ is the finesse of the cavity for the signal field. There is an enhancement in the optical depth by the effective number of round trips of the cavity $\mathcal{F}_s/\pi$ and an enhancement in the Raman coupling strength, C_s, by the same factor. The control field is also enhanced by the cavity and therefore there is a reduction in the control pulse energy required for the memory interaction by a factor $\sim \mathcal{F}_s^2\mathcal{F}_\Omega^2/\pi^4$, where $\mathcal{F}_\Omega$ is the finesse of the cavity for the control mode. The cavity enhances the memory interaction and a smaller control pulse energy is required to reach the same memory efficiency.

2. **Noise suppression**

   The signal to noise ratio (SNR) for the retrieved state from the Raman memory is given by:

$$\text{SNR} = N_{\text{in}} \left|\kappa\right|^2 \times \frac{E^2}{g(\zeta)} \times \left|\frac{C_s}{C_a x}\right|^2 \tag{4.39}$$

   where $g(\zeta) = (1 - \mathrm{e}^{-\zeta}E)/\zeta$. We see that the SNR increases with input photon number, N_{in}, and with the overlap between the optical fields and cavity modes, $|\kappa|^2$. E^2 quantifies the contribution to the retrieved signal from the beam splitter Green's function mapping the input signal $\sigma_{\text{in},0}$ to the output signal $\sigma_{\text{out},1}$, and $g(\zeta)$ captures the mapping from the anti-Stokes field to output signal via the four-wave mixing Green's function. The final term quantifies the ratio of the coupling strengths of the signal field, C_s and the suppressed anti-Stokes field, $C_a x$, and demonstrates that the cavity suppresses the unwanted anti-Stokes coupling by $1/\left|x\right|^2$. To maximise the suppression we tune to resonance for the signal field and anti-resonance for the anti-Stokes field ($\phi_s = 0, \phi_s = \pi$) and therefore get a noise suppression of $x = (1 - \mu_s)/(1 + \mu_a)$.

3. **Mode selectivity**

   If we look at the form of Equation 4.35 we can see that the mapping from the input signal, $\sigma_{\text{in},0}$, to the retrieved signal, $\sigma_{\text{out},1}$, is described by the Green's function $M_{\text{BS}}^{[1]}$. This can be rewritten as

$$\int_0^1 \mathrm{d}\epsilon' \left[M_{\text{BS}}^{[1]}(\epsilon, \epsilon')\sigma_{\text{in},0}(\epsilon')\right] = -\chi C_s^2 \mathrm{e}^{f(1+\epsilon)} \int_0^1 \mathrm{d}\epsilon' \left[\mathrm{e}^{-f\epsilon'}\sigma_{\text{in},0}(\epsilon')\right] \tag{4.40}$$

   and is separable in ϵ and ϵ'. This separability implies that the memory is perfectly mode selective: the function $\mathrm{e}^{-f\epsilon(t)}$ fully determines which input signal mode is addressed by the

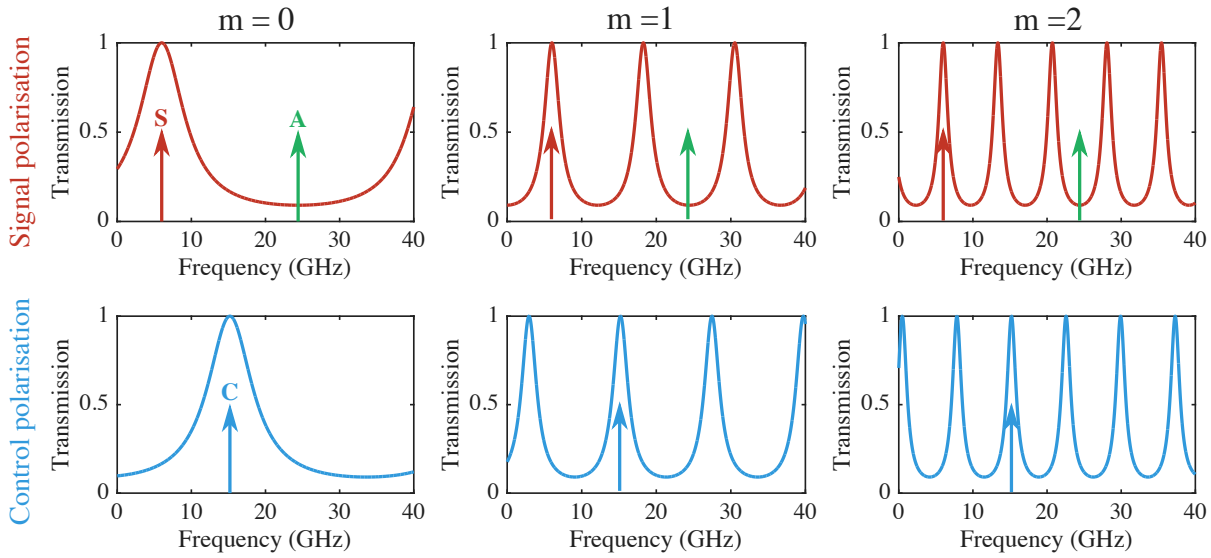


Figure 4.2: Transmission spectra for the three lowest order cavities, $m = 0, 1, 2$, where the free spectral range is given by $\mathrm{FSR}_m = 4\Delta_{\mathrm{hf}}/(2m+1)$. For the signal polarisation (red) the cavity is on resonance with the signal field (S) and anti-resonant for the anti-Stokes field (A). For the control polarisation (blue) the cavity is on resonance with the control field (C), and hence the "triple resonance" condition is satisfied. Note that these cavity spectra are for an "empty" cavity with no atoms, and do not account the for frequency dependent dispersion from the atomic line.

memory, and any input mode which is orthogonal to this will not couple to the memory at all. The stored mode can be arbitrarily chosen by shaping the control field and therefore the cavity-enhanced memory can act as a true temporal mode selective element [205]. This is a key point that we will come back to in Chapter 6.

This theoretical description demonstrates that operating the Raman memory inside a low-finesse optical cavity which is anti-resonant for the anti-Stokes field will enable suppression of the two-mode squeezer part of the Hamiltonian and enhancement of the beam splitter interaction, reducing the relative size of the noise process with respect to the desired memory process. The cavity will also enable higher memory efficiencies for a given control pulse energy and optical depth. In the next section I will outline the experimental infrastructure required to implement the cavity-enhanced Raman memory.

4.2 Experimental Implementation of a Cavity Memory

The first demonstration of the cavity-enhanced Raman memory was performed in 2015 by Saunders *et al.* [207] which was before I joined the research team. I will describe the experimental set-up of this *first-generation* cavity-memory, as much of the same infrastructure was used in the *second-generation* cavity that I built with D.J.S. and P.M.L.

To suppress four-wave mixing noise, we want to implement the memory in a low-finesse optical cavity which is on resonance for the signal and the control fields and anti-resonant for the anti-Stokes field. These three fields are each separated by the ground state splitting, Δ_{hf}, and therefore this "triple resonance" condition holds for a cavity with free spectral range

$\mathrm{FSR}_m = 4\Delta_{\mathrm{hf}}/(2m+1)$, or equivalently a round-trip optical length of $l_m = c/\mathrm{FSR}_m$, where m is an integer. This integer m indicates how many cavity modes there are between the signal and anti-Stokes frequencies, and we refer to m as the *order* of the cavity. The cavity frequency spectra for $m = 0, 1, 2$ are shown in Figure 4.2. For a fixed cavity bandwidth, the noise suppression factor, x, decreases with lower m. However, there is a large parameter space that determines the memory performance including the order of the cavity, m, reflectivity of the in-coupler, r_{in}, temperature of the caesium vapour, T, detuning from resonance, Δ, Rabi frequency of the control field, Ω, and the buffer gas pressure, p_{buff}. These parameters affect the linear absorption at the signal frequency, the finesse of the cavity, and the diffusion rate out of the interaction region, which in turn determine the memory efficiency, lifetime, bandwidth, and noise suppression. This large parameter space has been explored in detail by J.D.H.M. in [190]. For an initial demonstration, a second-order cavity was chosen with $\mathrm{FSR}_2 = 7.36$ GHz and $l_2 = 40.8$ mm which offers good memory efficiency and noise suppression with experimentally achievable parameters.

We choose a ring cavity geometry, rather that a linear one, to avoid generating a standing wave in the cavity. The nodes in the field would significantly reduce the memory efficiency and lifetime as atoms diffuse out of the small interaction region. The initial cavity design has two exit ports, one in transmission and one in reflection, as shown in Figure 4.3. This was chosen to ensure critical coupling of the incident fields into the cavity, and to allow a more straightforward experimental set-up with two means of optical access into the cavity. However, the retrieved signal can exit from either port of the cavity and therefore there is a maximum achievable efficiency of 50%.

The experimental set-up for the cavity-enhanced memory is shown in Figure 4.3. The pulse preparation of the signal and control fields is the same as described in Section 3.1.2. The signal and control pulses are selected from the Ti:Sapph pulse train using a Pockels cell, and are overlapped spatially and temporally before being aligned into the cavity. The stability condition of this cavity requires a beam waist of 70 μm located between the two flat mirrors [208], and a lens of focal length 400 mm is used to focus the signal and control beams into the cavity. The beam passes twice through a caesium vapour cell on each round-trip of the cavity and the vapour cell, of length 12 mm and diameter 10 mm, has anti-reflection coatings on the external windows to maximise the round-trip transmission of the optical fields. The vapour cell contains caesium and 10 Torr of neon buffer gas, and is heated to approximately 70°C using ohmic wire. The caesium ensemble is initialised in the ground state $|1\rangle$ via a counter-propagating optical pumping beam, as in Section 3.1.1. To achieve the required resonance conditions, a birefringent element is introduced to the cavity to allow fine control over the relative phase accumulation of the orthogonally polarised signal and control fields. The signal field is brought onto resonance with a cavity mode by slightly adjusting the temperature of the caesium vapour around the set point of 70°C, since the atomic dispersion changes the position of the cavity resonances. The birefringent element consists of 200 μm of $LiNbO_3$ and is thermally tuned, independently of the caesium vapour cell temperature, to bring the control field onto resonance with the orthogonally polarised cavity modes. An active feedback system was built by J.H.D.M. to independently stabilise the temperature of the caesium vapour and the $LiNbO_3$.

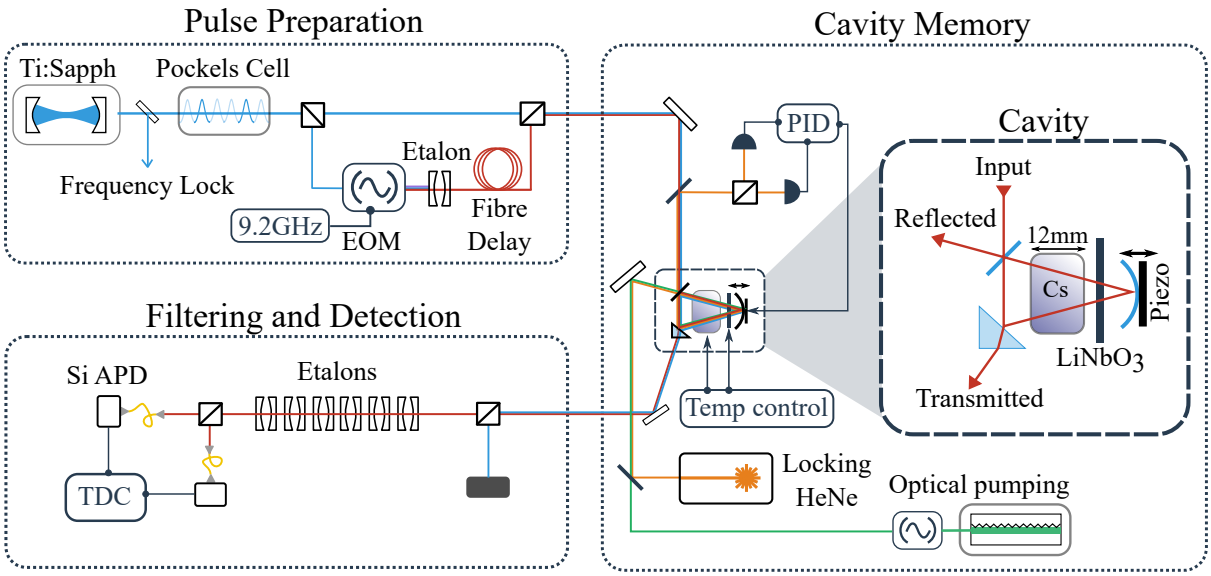


Figure 4.3: Experimental set-up for the cavity-enhanced Raman memory. The signal and control preparation, and the filtering and detection of the signal field are the same as in Chapter 3. The signal and control fields are focused into the cavity, and the output is detected through the transmitted mode. The caesium vapour cell and the birefringent $LiNbO_3$ element are temperature controlled to achieve the necessary resonance conditions. The cavity length is stabilised using a locking signal obtained from the polarisation rotation of a helium-neon (HeNe) laser which feeds back on the piezo controlled curved cavity mirror. The optical pumping beam is counter-propagating with the signal and control fields, and is switched off during the memory interaction using an EOM.

The length of the cavity is coarsely adjusted to be $L = l_2 = 40.8$ mm by adjusting the curved cavity mirror which is mounted on a translation stage. However, the cavity length needs to be stable to within the wavelength of the optical fields to maintain the resonance conditions. A locking signal is derived from a counter-propagating helium-neon (HeNe) laser using the Hänsch-Couillaud method. The HeNe light is prepared in diagonal polarisation with respect to the cavity, and the horizontal (H) and vertical (V) components acquire a different phase as they traverse the cavity due to the birefringence. The transmitted HeNe is separated into H and V using a polarising beam splitter (PBS), and the difference between these components is measured on two photodiodes. This difference signal is monitored and an Ardiuno is used to perform PID logic and feed back on the piezo-controlled cavity mirror to stabilise the cavity length. This mechanical stabilisation mechanism was developed by J.H.D.M, and allowed the cavity to be locked on resonance for timescales on the order of several hours.

The frequency response of the cavity is shown in Figure 4.4, where the transmission of a narrowband laser through the cavity is measured as it scans in frequency across the atomic resonances and the frequencies of the three optical fields. We can see that the signal field (detuned by $\Delta = 6$ GHz from the $6^2S_{1/2}\,|F = 3\rangle \rightarrow 6^2P_{3/2}$ transition) is on resonance with a cavity mode. The control field ($\Delta = 15.2$ GHz) is orthogonally polarised to the signal field, and has been brought onto resonance with the cavity modes using the birefringent element. The anti-Stokes field ($\Delta = 24.4$ GHz) has the same polarisation as the signal field, and is anti-resonant with the cavity as the length has been stabilised to $L = l_2$, and the cavity therefore

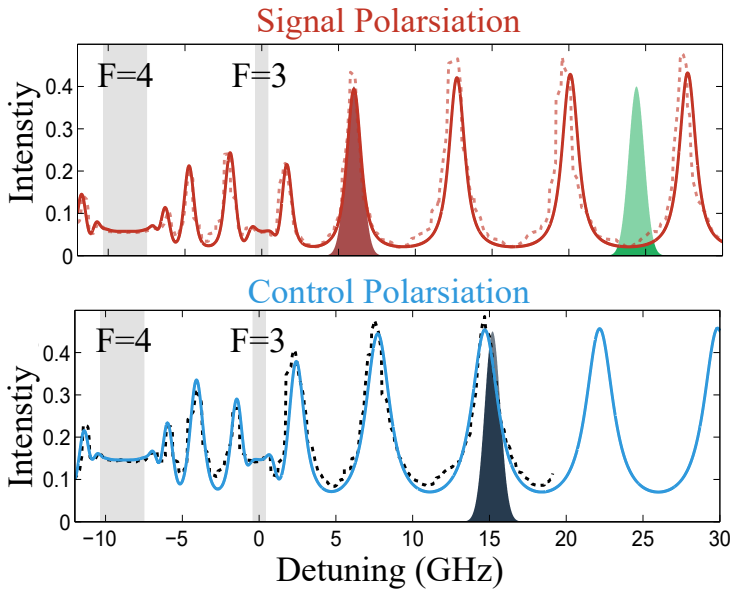


Figure 4.4: The transmission function of the cavity in the signal (upper) and control (lower) polarisations. The dashed lines are the measured values, and the solid lines are a fit to the data using a model developed by J.H.D.M. We see that the signal (red) and control (blue) fields are on resonance with the cavity, and the anti-Stokes field (green) is anti-resonant. The regions of absorption from the atomic transitions are shaded in grey. This figure is adapted with permission from [207].

satisfies the "triple-resonance" condition. The single-trip amplitude transmission of the signal and anti-Stokes field was measured to be $\mu_s = \mu_a = 0.6$, and therefore Equation 4.34 predicts a noise suppression factor of $x = (1 - \mu_s)/(1 + \mu_a) = 0.25$.

The filtering and detection of the retrieved signal from the memory are the same as in Section 3.1.3. The orthogonally polarised control field is separated from the signal using a calcite beam displacer. The signal is then transmitted through a series of Fabry-Pérot etalons to suppress the residual control field, before being coupled into a single-mode fibre and detected using a silicon APD.

The Raman memory was demonstrated in this cavity by storing and retrieving weak coherent states, and a memory efficiency of $\eta^{\text{cav1}} = (9.5 \pm 0.5)\%$ and a lifetime of $\tau^{\text{cav1}} = 95 \pm 7$ ns were achieved [207]. The number of noise photons per pulse was measured to be $N^{\text{cav1}}_{\text{noise}} = 0.015 \pm 0.002$ which gives a $\mu_1^{\text{cav1}} = N_{\text{noise}}/\eta = 0.17 \pm 0.02$; a factor of three improvement over the free-space Raman memory. We compare this to the cavity theory via Equations 4.38 and 4.39 which predict $N_{\text{noise}} = 0.005$ and $\mu_1 = 0.05$. This measured noise level is higher than predicted which could be due to another source of noise other than four-wave mixing. Nevertheless, this demonstrates the proof-of-principle of suppressing four-wave mixing noise via a cavity.

4.3 Second-Generation Cavity Memory

I have investigated how to improve the memory efficiency and noise suppression in the cavity-enhanced memory, to the level where the storage of single photons will be possible.

The original cavity design described above in Section 4.2 had two output ports, one in

reflection and one in transmission. The retrieved memory signal was measured in transmission only and this automatically reduced the memory efficiency by a factor of two, since the retrieved field could exit either port of the cavity. Furthermore the round-trip amplitude transmission of the signal was 0.6, and therefore the accumulated loss of the signal was very high.

We make some careful distinctions at this point to correctly account for this intra-cavity loss. We define the *internal* memory efficiency, η_{int}, as the ratio of the retrieved signal to the amount of signal that is transmitted through the cavity when no control pulse is present, which is the memory efficiency we directly measure on the detectors. However, the amount of signal detected when the memory is turned off is significantly lower than the input signal due to intra-cavity loss. We therefore also consider the *external* memory efficiency, η_{ext}, as the ratio of the retrieved signal to the input signal before the memory, which accounts for the internal efficiency of the memory as well as loss inside the cavity.

The intra-cavity loss in the first-generation cavity memory reduced the external memory efficiency to significantly less than 1%, or equivalently increased the external μ_1 to be far greater than 1. This meant that it was not possible to store single photons in this cavity demonstration. To enable storage of single photons in the cavity-enhanced Raman memory I investigated how to increase the memory efficiency, minimise intra-cavity loss, and maximise the noise suppression to achieve a lower external μ_1.

The second-generation cavity consists of one curved and one flat mirror, both with reflectivity $r \gtrsim 0.99$, and a flat mirror with $r \sim 0.8$ which acts as both the in- and out-coupler. This leads to a one-port ring cavity and should allow for greater memory efficiencies. The higher reflectivity also increases the finesse of the cavity, which will enable additional noise suppression. Great care was taken to minimise the intra-cavity loss by reducing reflections off each optical component, and the round-trip intensity transmission of the signal was measured to be 93%.

The triple resonance condition for the second-generation cavity is achieved using the same procedure as above, and the frequency response of the cavity is shown in Figure 4.5. As the cavity now has only one output mode, we measure the frequency response in "reflection" rather than in "transmission": when the scanning narrowband laser is off-resonant with the cavity, the light is reflected off the in-coupling mirror into the output mode and we see a maximum in the collected light. The cavity modes now appear as minima in the reflected light, rather than peaks in the transmitted light.

The optical pumping beam is counter-propagating with the memory fields, and is aligned into the cavity via the output spatial mode of the signal and control. The pumping light is combined with the signal and control beams using a dichroic filter which transmits the pumping light (895 nm, resonant with the D_1 transition), and reflects the signal and control modes (852 nm). Figure 4.5 shows the difference in the cavity modes when the optical pumping beam is switched on and off. We can see the significant shift in the cavity response due to dispersion from the atomic ensemble, even up to 30 GHz away from the atomic transition. This demonstrates that it is critical to check the triple resonance condition whilst the ensemble is optically pumped into the initial state for memory operation.

To streamline the experimental set-up we now derive the locking signal for mechanical stabilisation of the cavity from the optical pumping beam, rather than a HeNe laser. The

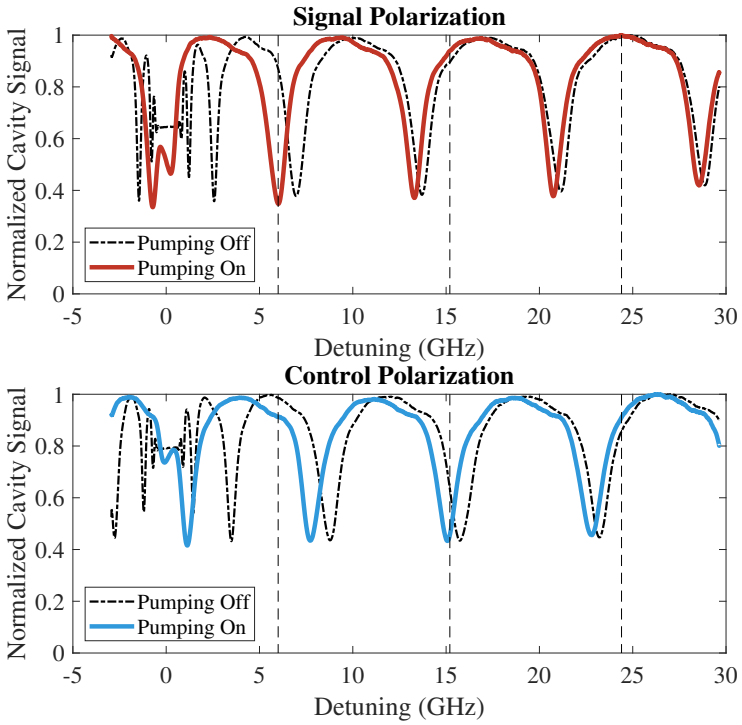


Figure 4.5: Frequency response of the second-generation cavity memory. The amount of light reflected into the output mode from the cavity in-coupling mirror is measured as a function of frequency using a narrowband scanning laser. The detuning is measured with respect to the $6^2S_{1/2}\,|F=3\rangle \rightarrow 6^2P_{3/2}$ transition. The dashed lines show the cavity response when the optical pumping light is turned off and the atomic population is equally distributed between the two ground states. The solid lines show the amount of reflected light for the signal (red) and control (blue) polarisation when the atomic ensemble is optically pumped into the $6^2S_{1/2}\,|F=4\rangle$ ground state, and only a small amount of population remains in the $6^2S_{1/2}\,|F=3\rangle$ state, and the cavity modes are significantly shifted due to the change in atomic dispersion.

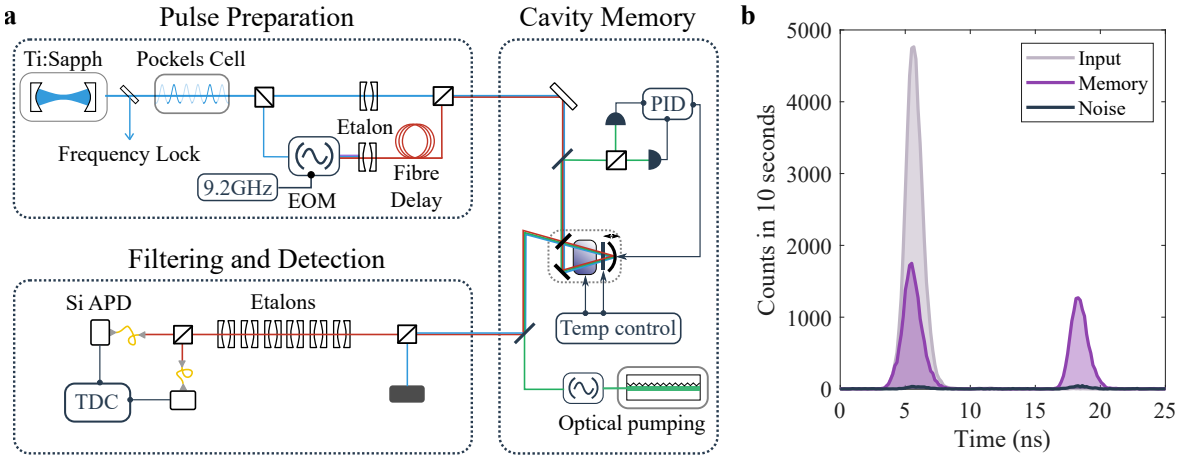


Figure 4.6: (a) Experimental set-up for the second-generation cavity-enhanced Raman memory. This is very similar to the previous cavity memory in Figure 4.3, except the cavity now has only one output mode. The counter-propagating optical pumping beam is combined with the output signal and control mode using a dichroic mirror. The locking signal to stabilise the cavity is obtained by measuring the polarisation rotation of the optical pumping beam through the cavity. (b) A typical arrival time histogram showing memory operation in the second-generation cavity memory.

cavity mirrors are coated to also have high reflectivity at the pumping wavelength, such that the optical pumping beam also experiences cavity modes. The polarisation rotation of the optical pumping light due to the cavity birefringence is measured in the same way as before, and a Hänsch-Couillaud locking signal is obtained. However, when the cavity is locked to be on resonance with the signal and control fields, it is not known where the cavity modes around 895 nm will be with respect to the frequency of the optical pumping beam. It so happens that when the cavity is locked to satisfy the triple resonance condition, the optical pumping beam is also close to resonance with the cavity and experiences a field enhancement. It is now sufficient to use 300 μW of light to achieve a high degree of spin-polarisation, compared to 3 mW in the non-cavity memory.

The experimental set-up for the second-generation cavity memory is shown in Figure 4.6(a). The generation of the signal and control pulses is the same as above, but we added a Fabry-Pérot etalon with a 18 GHz free-spectral range to remove any broadband noise arising from amplified spontaneous emission from the Ti:Sapph laser. This helps improve the noise floor, and has the added benefit of reducing the bandwidth of the control pulses so that they are better matched with the cavity modes, allowing for more injection of power into the cavity for a given pulse energy measured before the cavity. The transmission of this etalon at the control frequency is 70% and thus we still have sufficient pulse energy to drive the memory interaction. A typical memory trace for the second-generation cavity memory is shown in Figure 4.6(b), for a control pulse energy of 1 nJ.

We use the visibility of the cavity spectra to predict the intra-cavity loss for a signal pulse that is on-resonant with the cavity, and has a bandwidth that is 50% of the cavity linewidth. The cavity visibility is 0.43, from which we predict a ratio in the transmission for a pulse on-resonance with the cavity to an off-resonant pulse of around 0.20, as shown in Figure 4.7. We

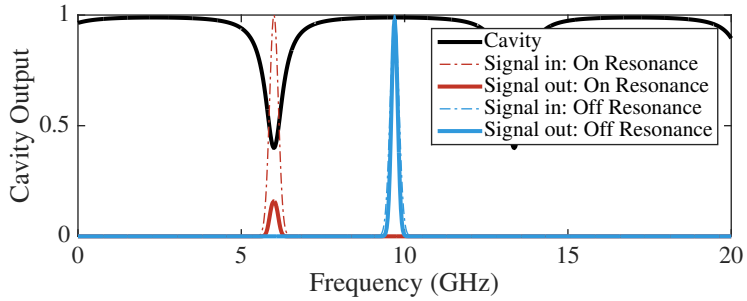


Figure 4.7: Simulation of the transmission of a pulse that is on-resonance (red) and off-resonance (blue) with an empty cavity spectrum (black). The pulse bandwidth is assumed to be half of the cavity bandwidth.

measure the transmission of an on-resonance pulse to be 10% which may indicate that photons are not optimally matched in bandwidth to the cavity. We therefore have $\eta_{\text{ext}} \approx 0.1 \times \eta_{\text{int}}$, which highlights the critical role of intra-cavity loss – even though the intensity transmission for a single round trip of the cavity is 93%, the accumulated loss of the signal is 90%.

The internal memory efficiency for the data in Figure 4.6(b) is $\eta_{\text{int}}^{\text{cav2}} = (25.7 \pm 0.3)\%$, which is significantly higher than in the first-generation cavity with $\eta_{\text{int}}^{\text{cav1}} = (9.5 \pm 0.5)\%$. The single-port cavity doubles the efficiency, and the remaining improvement is due to the increased overlap of the signal and control pulses to the cavity modes. We saw in Equation 4.38 that the efficiency increases with $|\kappa|^2$ which quantifies the overlap between the cavity modes and optical fields, and hence for a given pulse energy we have increased κ and improved the memory efficiency. The measured transmission from after the cavity to the detector is 6.0% (due to the improved performance of the etalons from Section 3.1.3), so we deduce that the input photon number is $N_{\text{in}} = 1.56$. We therefore measure $\mu_{1,\text{int}}^{\text{cav2}} = 0.056 \pm 0.003$ – a factor of 3 improvement over the previous cavity results.

To investigate the residual noise we apply multiple read-out pulses, each separated by 12.5 ns. We see in Figure 4.8(a) that there is retrieval of the signal over three subsequent read-out pulses. The retrieval efficiency here is 33%, and thus this imperfect conversion from the spin-wave into the retrieved signal in the first output bin leaves some residual spin-wave which can subsequently be read-out of the memory. Figure 4.8(b) shows the number of noise counts in each of these time bins, and we can see that it is approximately equal. We measure the number of noise photons per pulse in these four time bins as a function of control pulse energy, and the results are shown in Figure 4.8(c). We can see that the noise is approximately linear with control pulse energy, and that it is equal, within error, in all four time bins. If the noise was due to four-wave mixing we would expect it to increase quadratically with control energy, since this noise process involves two control photons. We would also expect the noise to increase with subsequent applications of the control pulse due to the generation and partial retrieval of a spin-wave with each pulse [208]. The fact that the noise is constant across multiple time bins, and increases linearly with control pulse energy indicates that it is not due to four-wave mixing. Other possible sources of noise include fluorescence from the excited state, leakage of the control field due to insufficient filtering, and spontaneous Stokes scattering due to imperfect optical pumping. Fluorescence noise will be distributed over a characteristic timescale of the

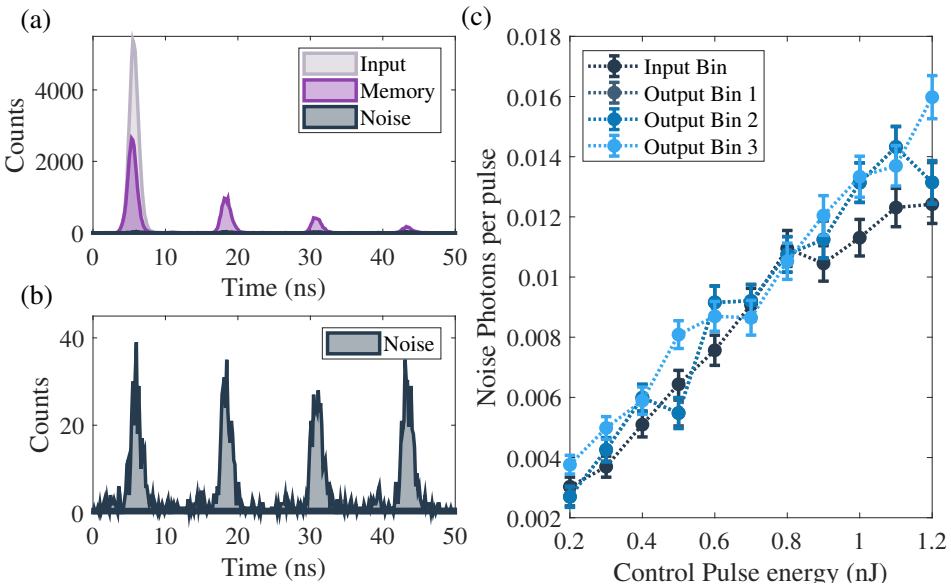


Figure 4.8: (a) Arrival time histogram showing the cavity-enhanced Raman memory with multiple read-out pulses. (b) The number of noise photons detected on the APDs when no input signal is present and four control pulses are applied. (c) Noise photons per pulse in the input time bin and three retrieval time bins as a function of control pulse energy.

excited state lifetime which is 30.5 ns, and therefore we can remove the majority of this noise by only integrating over the 330 ps pulse, but some fluorescence noise will remain in pulse time window. It is challenging to characterise what the exact source of the remaining noise is, and this is something we will come back to in Chapter 5. However, we note that whilst it is experimentally challenging to fully eliminate the strong control field or fluorescence noise, this is not an intrinsic problem and could be improved with better engineering using, for example, a custom made Bragg filter, or further improvement of the optical pumping efficiency.

This second-generation cavity has increased the memory efficiency by a factor of 3 to $\eta_{\text{int}}^{\text{cav2}} = (25.7 \pm 0.03)\%$. The noise metric, $\mu_{1,\text{int}}^{\text{cav2}} = 0.056 \pm 0.003$ is a factor of 3 lower than the previous cavity memory, $\mu_{1,\text{int}}^{\text{cav1}} = 0.17 \pm 0.02$, and a factor of 9 lower than the free-space Raman memory, $\mu_1^{\text{fs}} = 0.5$ [183], demonstrating that the cavity-enhanced Raman memory is effective in reducing four-wave mixing noise. However, we note that intra-cavity loss reduces the memory efficiency to $\eta_{\text{ext}} = (2.57 \pm 0.03)\%$, or equivalently increases the μ_1 to $\mu_{1,\text{ext}} = 0.56 \pm 0.03$. Whilst this is similar to the noise level in the free-space Raman memory, we expect that the relative contribution of various noise sources is now different due to the suppression of four-wave mixing noise, and therefore the resultant effect on a retrieved single photon will be different. By storing and recalling non-classical states we should be able to pin-point the various residual noise contributions in more detail, in a technique that was used successfully in [183], and something we will see again in Chapter 5.

The μ_1 parameter quantifies the retrieved noise-to-signal ratio for an input photon number of 1, and therefore if we have a single photon source with a probability of producing a single photon at the input of the memory of η_{h}, then the predicted SNR of the retrieved single photon is η_{h}/μ_1. If we can generate single photons before the cavity memory with a probability of

greater than 56% then we predict an output SNR greater than one. Whilst this is a challenging requirement, we nonetheless investigate the storage of single photons in this cavity-enhanced Raman memory.

4.4 Single Photon Source

To store single photons in the Raman memory we require a bright source of heralded single photons, which have a bandwidth of around 1 GHz and are detuned in frequency from the D_2 caesium transition at 852 nm by 15.2 GHz. The source we use is similar to that described in [208], but has been rebuilt by myself and B.B. to improve the performance.

4.4.1 SPDC Photon Source

We interface the cavity-enhanced Raman memory with a photon source based on spontaneous parametric down-conversion (SPDC). There is extensive literature on SPDC photon sources, for example see [209] for a review or [84, 86, 93, 94] for recent examples, and I will give a brief outline below.

SPDC is a second-order nonlinear process in which a pump photon (p) is converted into two, lower energy photons called the signal (s) and the idler (i). The down-conversion is known as type-I if the down-converted photons have the same polarisation, and type-II if they have orthogonal polarisation. SPDC occurs in materials with a non-zero second order susceptibility $\chi^{(2)}$, i.e. non-centrosymmetric crystals such as Potassium titanyl phosphate (KTP) and Beta barium borate (BBO) [210]. The process conserves both energy and momentum, therefore the frequencies and wavevectors of the down-converted photons satisfy:

$$\begin{aligned} \omega_\mathrm{p} &= \omega_\mathrm{s} + \omega_\mathrm{i} \\ \mathbf{k}_\mathrm{p} &= \mathbf{k}_\mathrm{s} + \mathbf{k}_\mathrm{i}. \end{aligned} \tag{4.41}$$

The energy conservation constraint means that the down-converted photons are correlated in frequency. Furthermore, only a small subset of frequencies of down-converted photons satisfy the momentum conservation condition (known as phase-matching) due to dispersion inside the medium. The phase-matching condition can be altered by changing the bandwidth of the pump laser or the temperature of the crystal, to control the frequency of the generated photons.

SPDC photons are always produced in pairs and the output state of a SPDC source is a two-mode squeezed vacuum:

$$|\Psi\rangle_\mathrm{SI} = \sqrt{1-\lambda^2}\sum_{n=0}^{\infty}\lambda^n\,|n\rangle_\mathrm{S}\,|n\rangle_\mathrm{I} \tag{4.42}$$

where $|n\rangle_{\mathrm{S/I}}$ is an n-photon Fock state in the signal/idler mode respectively, and λ is the squeezing parameter which is determined by properties of the nonlinear crystal and the power of the pump beam. By increasing the pumping power we can increase λ and therefore increase the single photon pair production rate, but this also increases probability of multi-photon Fock state generation. These multiphoton states increase the second-order autocorrelation, $g^{(2)}(0)$,

of the SPDC state, and therefore we need to keep the squeezing sufficiently low to suppress higher order events.

We use a photon source based on type-II spontaneous parametric down-conversion (SPDC) in a periodically-poled potassium-titanyl-phosphate (PPKTP) waveguide. We use a waveguide rather than a bulk $\chi^{(2)}$ crystal so that the emitted photons are in a single, well-defined spatial mode, to facilitate coupling into single-mode fibre, and to increase the brightness of the source for a given pump power. The PPKTP waveguide is pumped with pulses at 426nm to produce orthogonally polarised near-degenerate signal and idler photons at 852 nm. The 426 nm pulses are generated by passing the 852nm Ti:Sapph pulses through a separate, 2 mm long bulk PPKTP crystal which frequency-doubles the photons via second harmonic generation (SHG). The efficiency of this SHG process is low, and approximately 4 mW of blue light is generated from 1 W of Ti:Sapph light. The frequency-doubled light is separated from the remainder of the Ti:Sapph pulse using a dichroic mirror, leaving sufficient 852 nm light to act as the control pulse to drive the memory.

The PPKTP chip is 20 mm in length and consists of multiple waveguides of widths 2, 3, and 4 μm, and the blue light is focused into a single waveguide using a 11 mm aspheric lens. The transmission of the blue light through the waveguide is 35%, and the residual 426 nm light is removed from the generated 852 nm photons using an edge filter (Semrock). The signal and idler photons are orthogonally polarised and are separated on a polarising beam splitter before being coupled into two single-mode fibres (SMFs). To align these modes we use narrowband 852 nm light which is coupled to the fundamental mode of the waveguide, and aligned into the SMFs. By careful choice of focusing optics to map the spatial mode of the waveguide output to that of the SMFs, we were able to achieve a coupling efficiency of 70% for the signal mode and 68% for the idler mode. The output PDC spectrum was characterised using an Andor spectrometer and it was found that the 3 μm guides have near-degenerate emission of the signal and idler photons at 852 nm. The crystal was temperature-tuned to adjust the phase-matching condition, and at approximately 70°C the emission becomes degenerate, as shown in Figure 4.9.

4.4.2 Photon Source Filtering

The phase-matching condition of this source is very broad so the generated photons have a bandwidth of approximately 1 THz. This is several orders of magnitude broader that the memory bandwidth and therefore we spectrally filter the idler photon to project the signal photon onto a state that is centred at the memory signal frequency, and has a bandwidth comparable to the control field and cavity modes. The idler beam is passed through two Fabry-Pérot etalons, one with a FSR of 18 GHz and one with an FSR of 103 GHz, giving an effective FSR of 1 THz. These etalons are temperature-tuned to the idler frequency in the same way as described in Section 3.1.3, which significantly improves the transmission. The beam then diffracts off a holographic volume Bragg grating with a width of $\sim$ 100 GHz, and the resultant width of this filtering set-up is 1 GHz. The idler photons are then coupled into a multi-mode fibre and detected using an APD, and a detection of a photon heralds the existence of a signal photon. The total transmission of the idler photons from the output of the waveguide to the herald detector is 22%. The idler detection rate after filtering is approximately 25 kHz and the

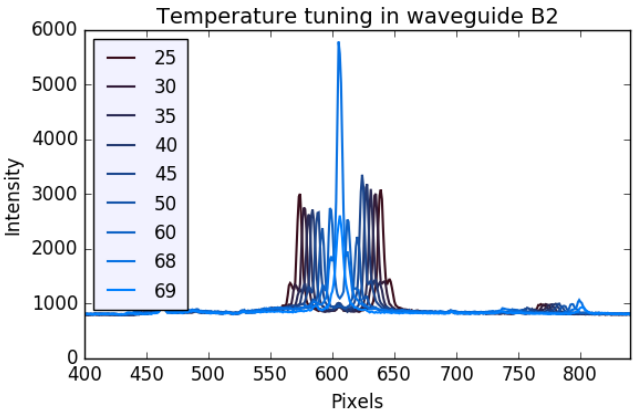


Figure 4.9: The measured PDC spectrum of one of the 3μm PPKTP waveguides as a function of temperature, measured using an Andor spectrometer. At room temperature there are two distinct peaks corresponding to the signal and idler frequencies. As the waveguide is heated the peaks change position due to the different phase-matching condition, and at $\sim 70°$C the signal and idler photons are degenerate.

single photon purity is $g^{(2)} \sim 0.02$ [89, 123].

4.4.3 Heralding Efficiency

We want to operate the memory in a feed-forward configuration, such that the detection of an idler photon triggers the memory operation, and this is described in more detail below. The heralding efficiency, η_{h}, is defined as the probability that we have a signal photon just before our memory, given that we detect an idler photon. We discussed above in Section 4.3 that the signal to noise ratio of a retrieved single photon from the memory is η_{h}/μ_1 and therefore we want to increase the heralding efficiency as much as possible.

To increase the heralding efficiency we need to maximise the transmission of the signal photons from the SPDC source to the memory. By carefully matching the spatial mode of the single photons generated by the source to that of the single-mode fibre we were able to reach a fibre coupling efficiency of 70%. We also streamlined the experimental set-up and minimised the number of optical components between the photon source and the memory, achieving a total transmission from the SPDC source to the memory (including the fibre coupling) of 65%. To measure the heralding efficiency, we couple the signal photons into a single-mode fibre just before the cavity memory and measure the photons using an APD. The idler photons are detected on a second APD and we use a time tagger to measure the photon count rates of the signal, c_{s}, and the idler, c_{i}, as well as the coincidences between the two detectors, $c_{\mathrm{s,i}}$. By correcting for the coupling efficiency of the fibre just before the memory, η_{fib}, and the efficiency of the detector $\eta_{\mathrm{D}} = 35\%$, we measure the heralding efficiency as

$$\eta_{\mathrm{h}} = \frac{c_{\mathrm{s,i}}}{c_{\mathrm{i}}\ \eta_{\mathrm{fib}}\ \eta_{\mathrm{D}}} \sim 30\%^{1}. \tag{4.43}$$

[1] In certain operational conditions the heralding efficiency of the source was measured to be as high as 55%. However, the stringent requirements on the filtering of the idler photons to project the signal photons onto a state centred at the signal frequency with the correct bandwidth reduces the heralding efficiency. This could be improved by optimising the filtering of the idler mode.

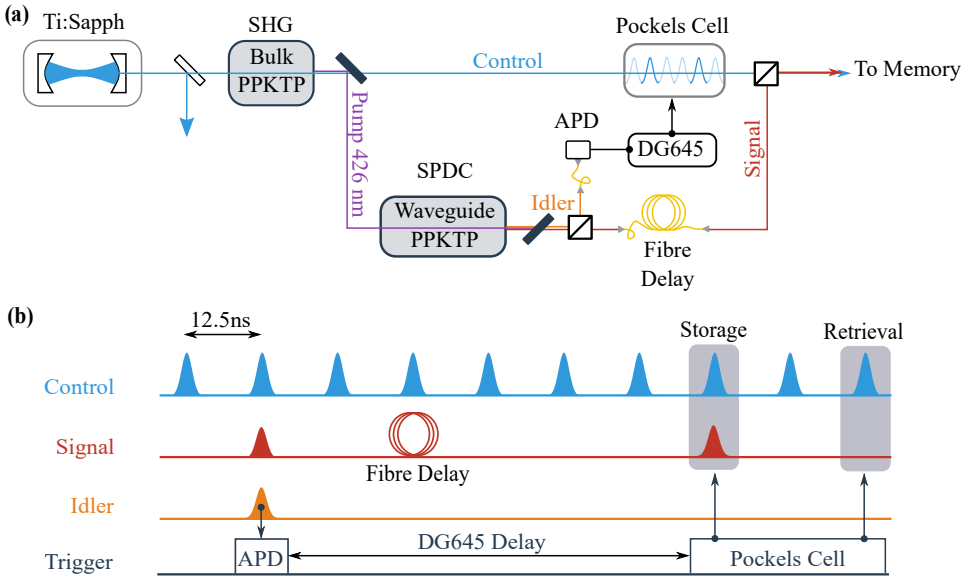


Figure 4.10: (a) A schematic showing feed-forward operation for storage of single photons in the memory. When a photon is detected in the idler arm, this triggers the Pockels cell to pick two pulses. The signal photon is delayed in an 80m single-mode fibre (SMF). The timing of the trigger is set such that the first control pulse is temporally overlapped with the signal photon. (b) The pulse sequence for feed-forward operation.

Given the μ_1 of the cavity-enhanced Raman memory is $\mu_{1,\mathrm{ext}}^{\mathrm{cav}} = 0.56 \pm 0.03$, then we predict a SNR of $\eta_{\mathrm{h}}/\mu_1 \sim 0.45$.

4.4.4 Feed Forward

We use an idler detection event to herald the production of a signal photon before the memory and to indicate that we should attempt storage. This is achieved by using the idler APD as a trigger for a digital delay generator (Stanford Research Systems DG645) which in turn triggers the Pockels cell. The Pockels cell then picks two control pulses, separated by the desired storage time of the memory. The signal photon is coupled into a 80 m long SMF to delay it by approximately 400 ns with respect to the control field, to ensure that there is sufficient time for the feed-forward electronics to trigger the Pockels cell. These feed-forward electronics were implemented by B.B. and D.J.S, and a schematic for the feed forward operation is shown in Figure 4.10.

It is crucial that the read-in control pulse arrives at the memory at exactly the same time as the signal photon. To match the path lengths of the signal and control we replaced the dichroic filter that separates the blue light generated via SHG from the remainder of the Ti:Sapph light with a beam splitter, and sent half of the 852 nm light through the PDC waveguide and into the long signal fibre and the other half into the control fibre. We recombined these fields just before the memory, and prepared both fields with the same amplitude and diagonal polarisation with respect to a polarising beam splitter. We observed the fluctuations at the output of the PBS, and adjusted the relative path length difference until we saw interference fringes with visibility $\gtrsim 90\%$. This implies that the two pulses are well overlapped temporally, and the intensity fluctuations are due to slight changes in the path length difference which causes the interference

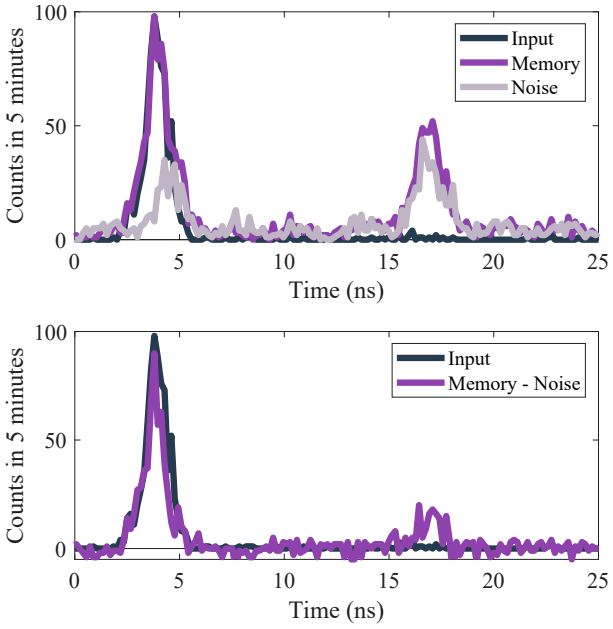


Figure 4.11: Memory trace for the storage of heralded single photons in the cavity-enhanced Raman memory. The noise is significantly higher than the signal in the retrieval bin.

to shift from constructive to destructive. The detection of an idler photon will therefore trigger the Pockels cell to pick two control pulses, the first of which will arrive at the memory at the same time as the delayed signal photon.

4.5 Interfacing the Cavity-Enhanced Raman Memory with a Single Photon Source

To test the noise suppression in the cavity-enhanced Raman memory I interfaced the second-generation cavity-enhanced Raman memory outlined in Section 4.3 with the heralded SPDC single-photon source described above in Section 4.4.

Figure 4.11 shows the storage of heralded single photons in the cavity-enhanced Raman memory. The measured memory efficiency here is $\eta = (16.6 \pm 0.7)\%$, which we note is lower than for the storage of weak coherent states in this cavity memory where $\eta = (25.7 \pm 0.3)\%$. This discrepancy is due to the difference in bandwidth between the single photons and the control pulse. This could be improved by altering the filtering of the idler photons to project the signal photons into a state that better matches the bandwidth of the control pulses. The drop in efficiency increases the effective μ_1 from 0.56 to 0.86, and therefore we would expect a signal to noise ratio of $\eta_h/\mu_1 = 0.30/0.86 = 0.35$. The signal to noise ratio of the output is measured here to be 0.3, which agrees fairly well with our prediction. The signal to noise ratio of retrieval is significantly less than one, and the retrieved signal is dominated by noise. The count rate of the photons after the memory was too low to be able to measure the second-order autocorrelation or perform any further noise investigations on the retrieved signal.

Whilst successful in storing single photons in the cavity-enhanced Raman memory, intra-cavity loss reduces the overall efficiency of the memory to a level where non-classical recall of

single photons is not possible, despite the significant reduction in noise. Reducing intra-cavity loss has therefore been highlighted as the key remaining challenge to enable quantum-level storage in the cavity-enhanced Raman memory. One method to address this problem would be to build a cavity where the optical elements are inside a caesium-filled vacuum cell, such that there are no optical surfaces inside the cavity. This is a challenging engineering problem since caesium is highly corrosive and therefore would gradually destroy any mirror coatings, and fine tuning the length of a cavity within a monolithic caesium cell is challenging. However, significant advances have been made in fabricating chip-scale caesium cells for atomic clocks and this may be the route forwards for a third-generation cavity memory. We are currently investigating other methods for noise suppression in Raman memories, but with additional engineering the cavity-enhanced Raman memory could suppress four-wave mixing to low enough levels for quantum storage.

4.6 Summary

In this chapter I have investigated using a low-finesse optical cavity to suppress unwanted four-wave mixing noise in the Raman memory, and the key results are outlined below.

- In Section 4.1 I outlined the theory behind the cavity-enhanced Raman memory, and presented the key theoretical results for the enhancement in memory efficiency, reduction in the noise, and increase in single mode capacity from [205].
- In Section 4.2 I described the experimental infrastructure for the cavity-enhanced Raman memory that was first demonstrated by Saunders *et al.* [207].
- In Section 4.3 I presented the next generation of the cavity memory that I have built, and demonstrated an improvement in the efficiency and the signal to noise ratio by a factor of three. This demonstrates that the cavity-enhanced Raman memory successfully suppresses four-wave mixing noise, and improves the SNR by a factor of 9 compared to the free-space Raman memory.
- In Section 4.4 I described the photon source that was built to interface with the cavity-enhanced memory. This photon source is based on spontaneous parametric down conversion, and generates heralded single photons with a 1 GHz bandwidth at the signal frequency, and a heralding efficiency of $\eta_{\mathrm{h}} = 30\%$.
- In Section 4.5 I discussed how I interfaced the second-generation cavity-enhanced Raman memory outlined in Section 4.3 with the heralded SPDC single-photon source described in Section 4.4. I observed the storage and retrieval of heralded single photons in the memory with an efficiency of $\eta = (16.6 \pm 0.7)\%$, but the signal to noise ratio of the output was less than one. I described how we have highlighted the key remaining challenge of the cavity-enhanced Raman memory as intra-cavity loss, which reduces the effective memory efficiency by a factor of 10.

We have investigated implementing the Raman memory in a low-finesse optical to suppress four-wave mixing noise, and have measured a factor of 9 improvement in the signal to noise

ratio. However, the efficiency is limited by intra-cavity loss and it is experimentally challenging to improve this. We are therefore investigating new, technically-simple methods of noise suppression which I describe in the following chapter.

Chapter 5:

Built-in Noise Suppression in a Raman Quantum Memory

In this chapter I present the results of a technically simple noise suppression scheme based on atomic absorption. This project was proposed and overseen by P.M.L, and the numerical simulations were performed by myself with assistance from T.M.H. and J.N. The experiment was performed by myself with assistance from T.M.H. and P.M.L.

5.1 Built-in Noise Suppression

We discussed in Chapter 2 that the Raman memory suffers from four-wave mixing noise which prohibits operation at the quantum level. In Chapter 4 we investigated a cavity-enhanced Raman memory, where a low finesse optical cavity suppresses anti-Stokes scattering by reducing the density of states at the anti-Stokes frequency. Whilst successful in reducing four-wave mixing noise, this method suffers from intra-cavity loss which severely limits the memory efficiency. It remains a significant engineering challenge to build a cavity which satisfies the necessary resonance conditions with ultra-low intra-cavity loss.

In this chapter we investigate an alternative method to suppress four-wave mixing noise. In Section 2.3.2 we discussed that the potential methods to suppress unwanted anti-Stokes scattering include selection rules, phase-matching and altering the density of states. Here we propose and demonstrate a technically simple noise suppression method that is based on atomic absorption or, equivalently, the introduction of a dispersion feature to destroy the four-wave mixing *phase matching* condition.

In 2016 Romanov *et al.* proposed a new method for reducing four-wave mixing noise in quantum memories by introducing an absorption feature on resonance with the noise field [211, 212]. This was demonstrated for an EIT memory in rubidium which suffers from four-wave mixing noise due to unwanted Stokes scattering. In this scheme they take advantage of the

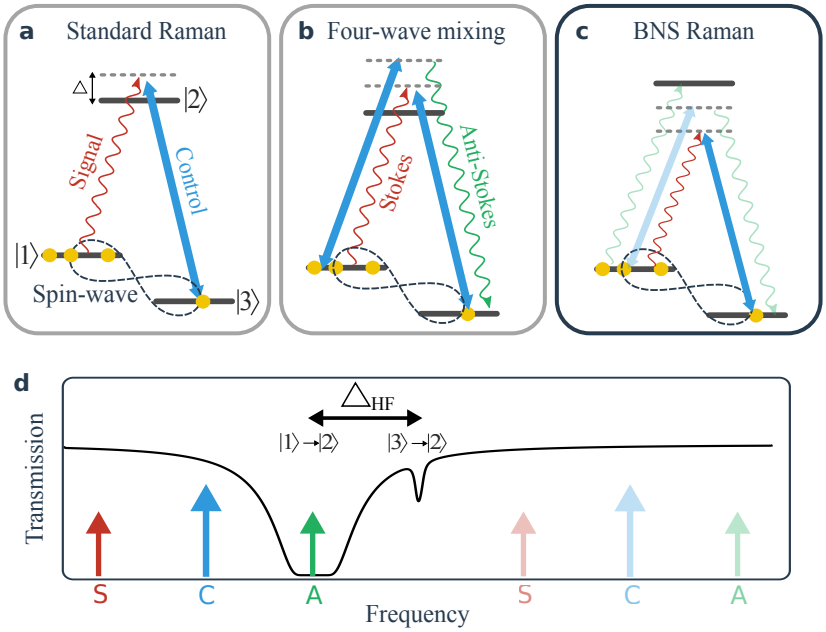


Figure 5.1: Raman Memory Schematic. (a) The Raman Memory protocol used a strong control field to drive a Raman transition and store an input signal field as a coherence across an atomic ensemble, or spin-wave. (b) Four-wave mixing noise arises when the control field couples to the populated state $|1\rangle$ and drives anti-Stokes scattering. (c) The "built-in noise suppression" (BNS) Raman memory operates at a specific detuning from resonance, $\Delta = -2\Delta_{\mathrm{hf}}$, such that the anti-Stokes field is on resonance with the populated $|1\rangle \rightarrow |2\rangle$ transition, and anti-Stokes scattering is strongly suppressed. (d) Simulation of the absorption spectrum of warm caesium vapour in a nitrogen buffer gas at 83°C, with 99.9% of the population in state $|1\rangle$. The arrows show the frequency of the signal, control and anti-Stokes fields for the BNS- (left) and standard- (right) Raman protocols.

two naturally occurring isotopes in rubidium and implement the memory in ^{87}Rb and use a strong control field to create a two-photon absorption feature with the Stokes field in ^{85}Rb. It was demonstrated that this scheme preserves the coherence in the EIT memory and suppresses four-wave mixing by 85%. However, since the absorption feature used to suppress the noise is via an off-resonant two-photon absorption, the suppression is not very strong.

In this chapter I present an adaptation of this protocol which is applicable to all atomic species. We modify the memory interaction by operating at a specific detuning from resonance such that the anti-Stokes frequency is resonant with the populated atomic transition of the atoms themselves, as illustrated in Figure 5.1. The anti-Stokes field therefore undergoes strong linear absorption and dispersion, which strongly suppresses anti-Stokes scattering and inhibits four-wave mixing noise. Simply by operating the memory at a specific detuning we utilise this built-in noise suppression (BNS) mechanism, achieving significant noise suppression without any detrimental effects on the memory efficiency or lifetime since the equations of motion that govern the memory interaction are unchanged. This method does not require any additional laser fields or cavity engineering and therefore holds promise as a route towards a technically simple, noise-free quantum memory.

5.2 Theory and Simulations

In this section I will outline the theory of this built-in noise suppression scheme and present the results of our numerical simulations of the system.

In Chapter 2 we derived the equations of motion for the signal field, S, and the anti-Stokes field, A, interacting with an ensemble of atoms via a strong control pulse with Rabi frequency Ω, and generating a spin-wave excitation B (Equations 2.47 – 2.49) which we repeat here for clarity:

$$(c\partial_z + \partial_t)\, S = \mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_s}B - \kappa_s S \tag{5.1}$$

$$(c\partial_z + \partial_t)\, A = \mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_a}B^\dagger - \kappa_a A \tag{5.2}$$

$$\begin{aligned}\partial_t B = &-\mathrm{i}\Omega^*\sqrt{\frac{d\gamma}{L}}\left[\frac{1-\alpha}{\Gamma_s} + \frac{\alpha}{\Gamma_s^*}\right]S + \mathrm{i}\Omega^*\sqrt{\frac{d\gamma}{L}}\left[\frac{1-\alpha}{\Gamma_a} + \frac{\alpha}{\Gamma_a^*}\right]A^\dagger \\ &- \left[\frac{1}{\Gamma_a^*} + \frac{1}{\Gamma_s}\right]|\Omega|^2 B.\end{aligned} \tag{5.3}$$

Here $\Gamma_{s,a} = \gamma + \mathrm{i}\Delta_{s,a}$ is the complex detuning of the signal and anti-Stokes fields, γ is the linewidth of the atomic transition (which can be broadened due to collisions with a buffer gas), and α is the proportion of atoms that have not been successfully initialised via optical pumping and remain in state $|3\rangle$. d is the optical depth of the ensemble, which is related to the on-resonance optical depth, d_0, by $d = d_0\gamma_N/\gamma$, where γ_N is the natural linewidth. The atomic absorption and dispersion are described by

$$\kappa_s = \frac{d\gamma c}{L}\left[\frac{(1-\alpha)}{\Gamma_s} + \frac{\alpha}{\Gamma_s^* + \mathrm{i}\Delta_{\mathrm{hf}}}\right] \tag{5.4}$$

$$\kappa_a = \frac{d\gamma c}{L}\left[\frac{(1-\alpha)}{\Gamma_a + \mathrm{i}\Delta_{\mathrm{hf}}} + \frac{\alpha}{\Gamma_a^*}\right], \tag{5.5}$$

where Δ_{hf} is the ground state splitting between hyperfine levels $|1\rangle$ and $|3\rangle$.

If we choose $\Delta_s = -2\Delta_{\mathrm{hf}}$, then the anti-Stokes field is on resonance with the $|1\rangle \rightarrow |2\rangle$ transition, since $\Delta_a = \Delta_s + \Delta_{\mathrm{hf}}$, and undergoes strong linear absorption. Equation 5.5 becomes (in the case of ideal pumping, $\alpha \rightarrow 0$), $\kappa_a \approx \frac{dc}{L}$, and by substituting this into Equation 5.2 we see that the absorption length for the anti-Stokes field in the atomic medium is on the order of $\sim L/d$, which for warm alkali vapours is typically sub 100 μm. As we shall see below, this prohibits anti-Stokes scattering and therefore four-wave mixing noise is significantly suppressed at this specific detuning $\Delta_s = -2\Delta_{\mathrm{hf}}$.

There are several ways we can consider the physical mechanism behind this noise suppression process. Firstly, the signal and anti-Stokes fields now experience vastly different propagation conditions, with the anti-Stokes field being strongly absorbed. The anti-Stokes field experiences a large amount of dispersion, and this destroys the phase-matching condition. By tailoring the frequency of the fields such that there is an absorption feature at the anti-Stokes frequency,

we can ensure that the four-wave mixing process is no longer phase-matched and strongly suppressed. Alternatively we consider that if an anti-Stokes photon is scattered it will quickly be absorbed over a characteristic length scale of L/d and the generated spin-wave excitation would therefore be localised. We can no longer generate a de-localised excitation over the entire atomic ensemble, as in the case of the desired memory interaction. The spatial confinement of the noisy spin-wave in this way prevents a coherent build up across the ensemble, and the resultant overlap with the memory spin-wave mode is greatly diminished. The absorption of an anti-Stokes photon acts as a measurement process which prevents the build up of coherence of the noise excitation and significantly limits the occupation of the noisy spin-wave, in a way that is akin to the dissipative quantum Zeno effect [213].

To test this absorptive noise suppression scheme we numerically solve the equations of motion around this specific detuning $\Delta_s = -2\Delta_{\mathrm{hf}}$. These simulations presented below have been performed by myself, with assistance from T.M.H. and J.N. We numerically solve Equations 5.1 – 5.3 using a combination of the Runge-Kutta method and Chebyshev iteration method in the following way:

- A grid in (z, t) space is used with (N_z, N_t) points. N_z must be sufficiently large to sample the short absorption length of the anti-Stokes field through the atomic medium, and N_t must satisfy $N_t > (c\,T_{\mathrm{sim}}/L)\,N_z^2$ where T_{sim} is the total simulation time, to ensure numerical stability.

- A matrix ρ combining S, A, B is initialised according to the boundary conditions at $t = 0$.

- A function f is created to perform derivatives: $\dot{\rho} = f(t, \rho)$. The Chebyshev method is used to perform the differential $\partial/\partial z$ on S and A.

- The Runge-Kutta method is used to perform the time derivative, by discretising ρ into N_t time steps and defining:

$$\begin{aligned} \rho_{n+1} &= \rho_n + \frac{dt}{6}(k_1 + 2k_2 + 2k_3 + k_4) \\ t_{n+1} &= t_n + dt \end{aligned}$$

  where

$$\begin{aligned} k_1 &= f(t_n, \rho_n) \\ k_2 &= f(t_n + \frac{dt}{2}, \rho_n + k_1\frac{dt}{2}) \\ k_3 &= f(t_n + \frac{dt}{2}, \rho_n + k_2\frac{dt}{2}) \\ k_4 &= f(t_n + dt, \rho_n + k_3 dt) \end{aligned}$$

- This results in a $(3 \times N_z \times N_t)$ array in (z, t) space for the three variables S, A, B.

We define the memory efficiency as the ratio of the energy of the retrieved signal to the input state:

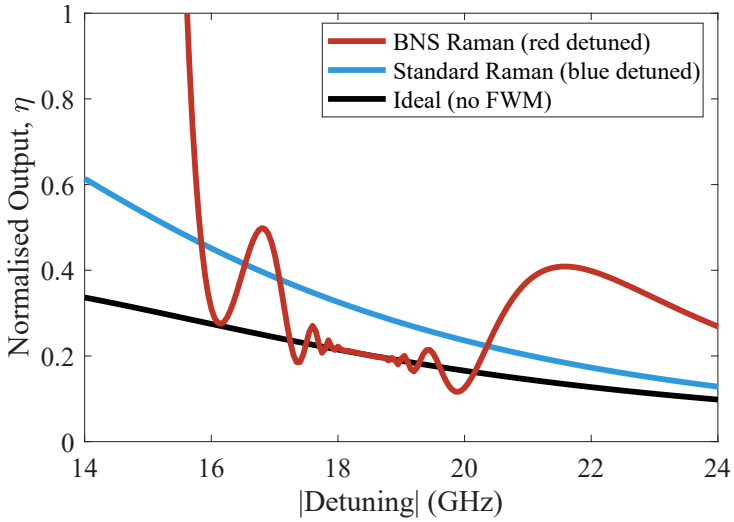


Figure 5.2: Numerical simulations for the efficiency of the Raman memory as a function of detuning for three configurations: blue-detuned from resonance in the standard Raman memory configuration (blue); red-detuned from resonance where the built-in noise suppression condition occurs at $\Delta = -2\Delta_{\mathrm{hf}} = -18.4$ GHz (red); and the ideal equations of motion with no four-wave mixing (black). The oscillations in efficiency are due to the phase-matching condition for four-wave mixing. At $\Delta = -2\Delta_{\mathrm{hf}}$ we see a significant suppression in gain due to four-wave mixing noise. These simulations were performed with $\gamma = \gamma_N = 5$ MHz, $d = 2 \times 10^3$ and integrated Rabi frequency $W = 21$ GHz.

$$\eta_{\mathrm{mem}} = \frac{N_{\mathrm{out}}}{N_{\mathrm{in}}} = \frac{\int_{-\infty}^{+\infty} \mathrm{d}t \left\langle S_{\mathrm{out}}^{\dagger}(t) S_{\mathrm{out}}(t) \right\rangle}{\int_{-\infty}^{+\infty} \mathrm{d}t \left\langle S_{\mathrm{in}}^{\dagger}(t) S_{\mathrm{in}}(t) \right\rangle}.$$

We consider three cases when simulating the memory efficiency numerically:

1. **BNS-Raman.** We consider a range of detunings around the absorption condition for built-in noise suppression $\Delta_s = -2\Delta_{\mathrm{hf}}$.

2. **Standard-Raman.** A similar range of detunings as in (1) is used for the blue-detuned case $\Delta_s = +2\Delta_{\mathrm{hf}}$.

3. **Ideal-Raman.** We artificially turn off the four-wave mixing process and calculate the memory efficiency for the ideal, red-detuned Raman memory equations.

Note that the efficiency for cases (1) and (2) can be greater than one due to four-wave mixing gain which amplifies the input signal, as we saw in Chapter 3.

For an initial demonstration, we perform these numerical simulations in an ensemble with no pressure broadening, ($\gamma = \gamma_N = 5.2$ MHz), and a relatively low optical depth of $d = 2 \times 10^3$. This allows us to explore how the system behaves in a parameter regime where the number of steps in z required to capture the anti-Stokes absorption is around $N_z = 20$, which is not prohibitively large. The results of these simulations are shown in Figure 5.2, and there are several key features that we draw our attention to. Firstly, when far away from the BNS condition, we see that the memory efficiency in case (1) is significantly higher compared to cases (2) and

(3). This is because, in this configuration, the coupling strength of the two-mode squeezing Hamiltonian, C_a is larger that the beam splitter coupling C_s since $\Gamma_a < \Gamma_s$, and therefore the four-wave mixing gain dominates. Secondly, as the detuning approaches the BNS condition $\Delta = -2\Delta_{\mathrm{hf}} = -18.4$ GHz, the anti-Stokes field undergoes significant atomic absorption, and the memory efficiency in case (1) tends towards that of the ideal case (3). At the exact BNS condition, the artificial increase observed in the memory efficiency due to four-wave mixing gain is suppressed by many orders of magnitude compared to the standard, blue-detuned memory of case (2).

In Figure 5.2 we see that there are oscillations in the memory efficiency either side of the absorption condition. To explore this in more detail we consider the phase-matching condition for the four-wave mixing process. The wavevectors of the signal, control and anti-Stokes fields are (for perfect optical pumping):

$$\begin{aligned} k_s &= \frac{\omega_s}{c} + \frac{d\gamma}{L}\left|\frac{1}{\gamma + \mathrm{i}\Delta}\right| \\ k_c &= \frac{\omega_c}{c} + \frac{d\gamma}{L}\left|\frac{1}{\gamma + \mathrm{i}(\Delta + \Delta_{\mathrm{hf}})}\right| \\ k_a &= \frac{\omega_a}{c} + \frac{d\gamma}{L}\left|\frac{1}{\gamma + \mathrm{i}(\Delta + 2\Delta_{\mathrm{hf}})}\right|, \end{aligned} \tag{5.6}$$

and the four-wave mixing phase mismatch is given by $\delta k = 2k_c - k_s - k_a$. When $\Delta = -2\Delta_{\mathrm{hf}}$ the anti-Stokes field undergoes significant dispersion and the second term in $k_a \to d/L \gg 1$. The phase mismatch, δk undergoes rapid oscillations modulo 2π, and the four-wave mixing process goes in and out of phase. This can be thought of as energy transfer back and forth between the four fields involved in four-wave mixing, resulting in a net transfer of energy between the signal field, the anti-Stokes field and the spin-wave. We see in Figure 5.3 that the oscillations in the memory efficiency match up well with the phase mismatch cycling through 2π. When $\sin(\delta k L) = 1$ we see a local maximum in the memory efficiency since the four-wave mixing process is optically phase matched. Similarly, when $\sin(\delta k L) = -1$ there is optimal destructive interference and there is a local minimum in efficiency since most of the energy is in the stored spin-wave. For the cases when $\sin(\delta k) = 0$, the four-wave mixing process is not phase-matched and the efficiency is the same as for the ideal equations of motion.

We change the optical depth of the simulations, and see in Figure 5.3 that the phase matching oscillations again match up with the variations we see in the efficiency. We also change the control pulse energy for a fixed optical depth and see that the frequency and position of the oscillations remain constant, but the amplitude increases with higher control energy as shown in Figure 5.4. The dispersion experienced by the anti-Stokes field is fixed by the optical depth, and therefore the phase mismatch δk does not change with control pulse energy. As the control pulse energy increases the four-wave mixing coupling strength increases, and therefore we see more gain in the system at the detunings when four-wave mixing is phase matched.

We now increase the resonant optical depth to more experimentally realistic values, and the results of these simulations for $d_0 = 3 \times 10^4$ (or a vapour temperature of 82°C) are shown in Figure 5.5(a). To fully capture the dynamics, the number of sampling points at $\Delta = -2\Delta_{\mathrm{hf}}$ has

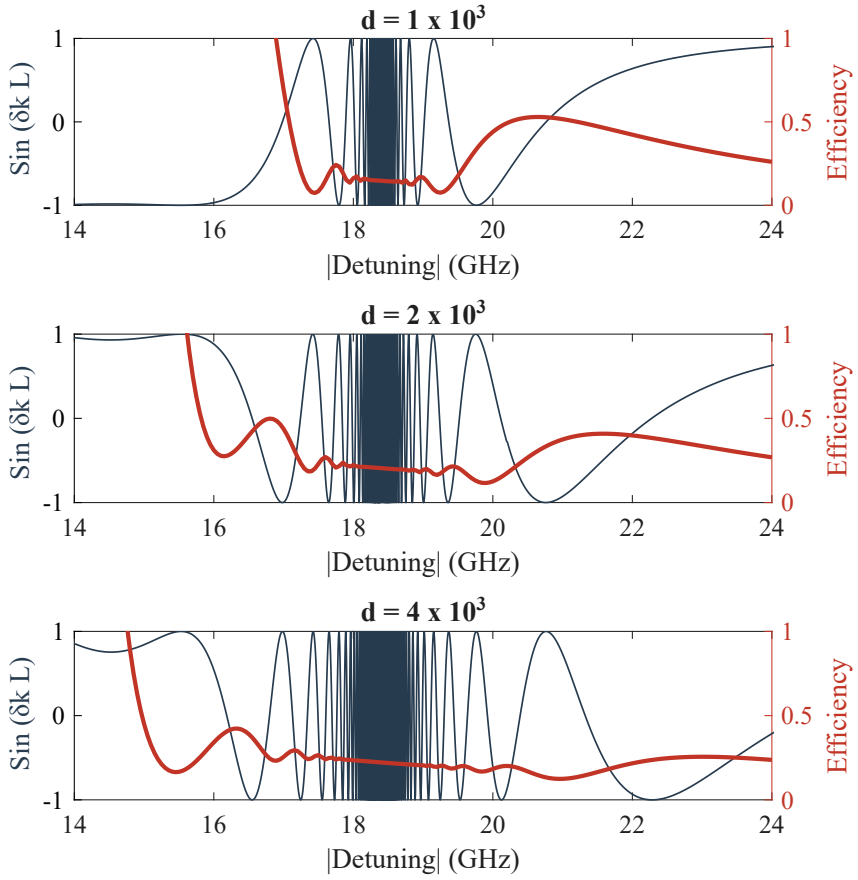


Figure 5.3: The phase matching condition and memory efficiency as a function of detuning from resonance for the different optical depths. The phase mismatch, δk, changes rapidly due to the dispersion of the anti-Stokes field and $\sin(L \times \delta k)$ oscillates around the resonance condition $\Delta = -2\Delta_{\mathrm{hf}}$. These oscillations match well with the variations in efficiency from the numerical simulations. The integrated Rabi frequency in these simulations is $\Omega = \{1/\sqrt{2}, 1, \sqrt{2}\} \times 21$ GHz for $d = \{1, 2, 4\} \times 10^3$ to keep the Raman coupling constant the same in all three cases. The linewidth is fixed at $\gamma = \gamma_N = 5.2$ MHz.

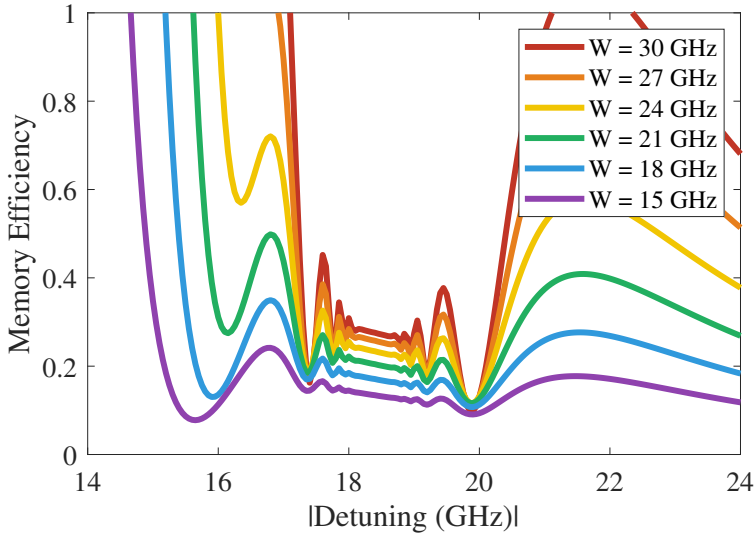


Figure 5.4: Memory efficiency as a function of detuning around the BNS-Raman condition for different integrated Rabi frequencies of the control, W. The optical depth is $d = 2 \times 10^3$ and the linewidth is $\gamma = \gamma_N = 5.2$ MHz.

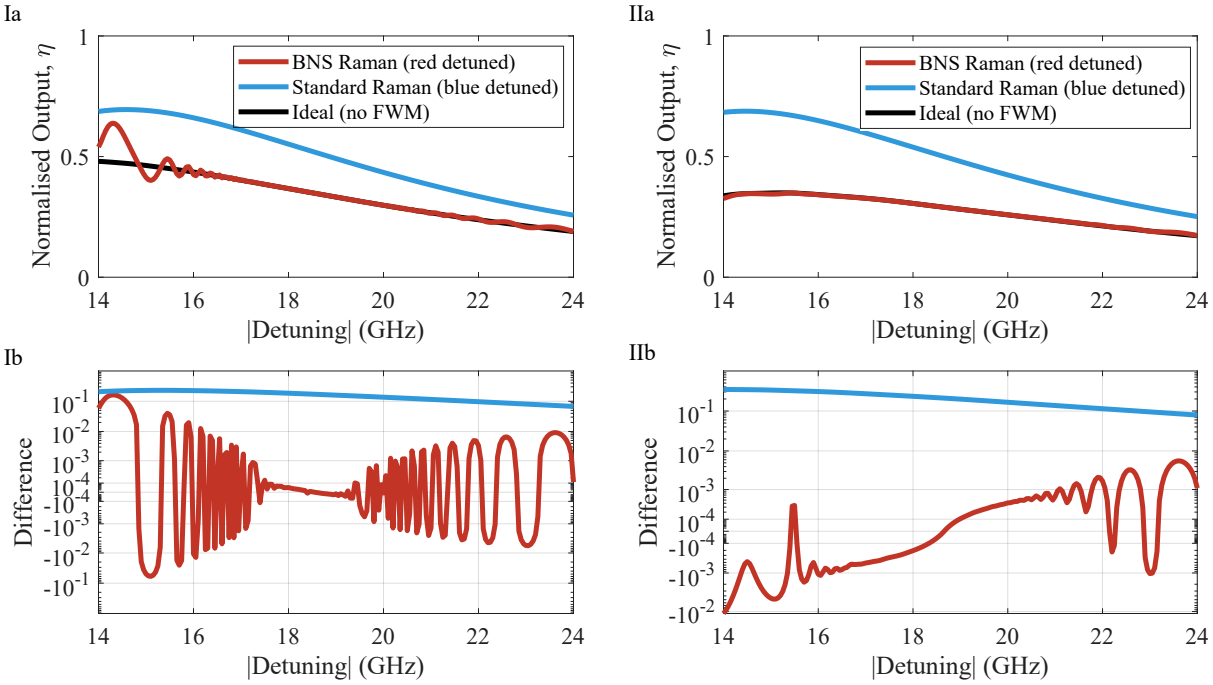


Figure 5.5: (a) Memory efficiency as a function of detuning for $d = 3 \times 10^4$, or a vapour temperature of 82°C, and (I) $\gamma = \gamma_N = 5$ MHz and (II) $\gamma = 95$ MHz. (b) The difference between the memory efficiency for the full equations of motion (including four-wave mixing) and the ideal equations (with no FWM) for the standard Raman memory (blue) and the Raman memory with built-in noise suppression (red).

to be increased to $N_z = 200$, $N_t = 600,000$, which slows down the simulations time significantly. We also include pressure broadening by increasing the linewidth to $\gamma = 95$ MHz, as shown in Figure 5.5(b). We can see that the region over which four-wave mixing noise is suppressed due to the BNS mechanism extends over several GHz. This is because the dispersion feature is very broad, and four-wave mixing is not phase-matched over a very wide range of frequencies. This demonstrates that this BNS mechanism is suitable for broadband memory protocols, where the noise field can itself have a bandwidth in excess of a GHz.

The lower plots in Figure 5.5 show the difference between the efficiency for the BNS-Raman (case 1) and the ideal-Raman (case 3). We can see that the relative difference between the BNS Raman memory and the ideal Raman memory is around four orders of magnitude lower than the standard Raman for more than a gigahertz around $\Delta = -2\Delta_{\mathrm{hf}}$. We therefore have strong confidence that this method will successfully reduce four-wave mixing noise even for broadband memories.

Equations 5.1 – 5.3 do not include all possible sources of noise, and hence these predictions give an upper bound on the performance of this scheme. We will explore other possible sources of noise and how these affect the retrieved signal below. Nevertheless, these simulations demonstrate the powerful efficacy of this method for the suppression of four-wave mixing noise.

5.3 Experimental Results

We investigate this built-in noise suppression scheme via the storage of weak coherent pulses in a Raman quantum memory in warm atomic caesium vapour. The experimental set-up is the

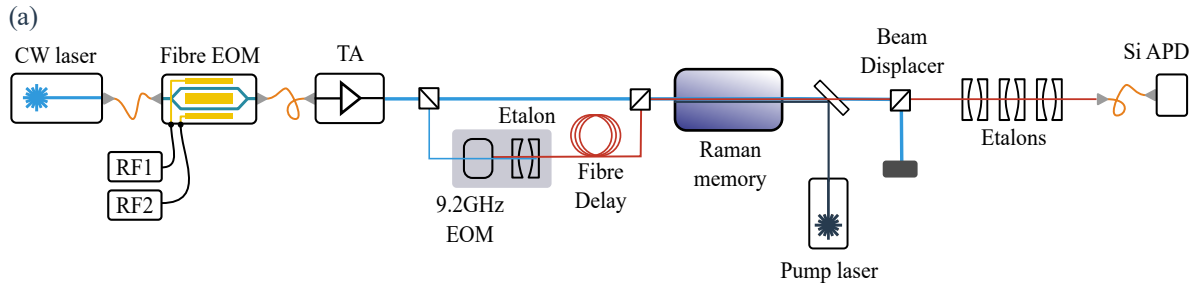


Figure 5.6: A schematic of the BNS-Raman memory experiment. Pulses are carved from a CW laser using a fibre-integrated electro-optic interferometer, and amplified using a tapered amplifier (TA). A small portion of the light is split off and passed through an EOM driven at 9.2 GHz to generate a sideband at the signal frequency, and the carrier frequency and other sideband are filtered away using an etalon. The signal is coupled into a long delay fibre, and then the signal and control beams are overlapped temporally and spatially inside a warm caesium vapour cell. The control field is filtered away using a series of polarisation and frequency filters. The caesium ensemble is prepared via a counter-propagating optical pumping beam. For more details, see Section 6.2.

same as in Chapter 3, except the signal and control pulses are now carved from a continuous wave generator using a Mach–Zehnder interferometer, rather than being derived from a pulsed Ti:Sapph laser. A schematic for the set-up is shown in Figure 5.6, and is described in detail below in Chapter 6. The signal and control pulses have a duration (full-width at half maximum of the intensity profile) of 10 ns, which is significantly longer than the pulse from the Ti:Sapph laser.

To test the absorptive noise suppression we operate the memory in two configurations: firstly at $\Delta_s = -2\Delta_{\mathrm{hf}}$ i.e. 18.4 GHz red detuned from the excited state, where the anti-Stokes field is on resonance with the atomic transition and therefore strongly absorbed (BNS-Raman), and secondly at 18.4 GHz blue detuned from the excited state, where the coupling strength of the Raman memory interaction is the same but there is no atomic suppression of the four-wave mixing noise (standard-Raman). The energy level diagrams for these regimes are shown in Figure 5.1.

5.3.1 Efficiency and Storage Time

We first measure the efficiency as a function of the control pulse energy in these two regimes. The results are shown in Figure 5.7, and the dashed lines are the predictions from our numerical simulations. There is strong agreement between the simulations and the experimental results which gives us confidence in our numerical model. We can see that the memory efficiency at low control pulse energies is similar in the two cases, but the efficiency in the BNS memory deviates at high control pulse energies. There are two reasons for this discrepancy at high coupling strengths. Firstly, we have suppressed the four-wave mixing process which otherwise artificially increases the memory efficiency due to gain at high coupling strengths. Secondly, the control pulse is closer to the populated state in the BNS configuration which induces a larger AC Stark shift that affects the two-photon resonance condition. We note that this is not a fundamental problem and can be counteracted with appropriate shaping of the temporal phase of the control pulses [205].

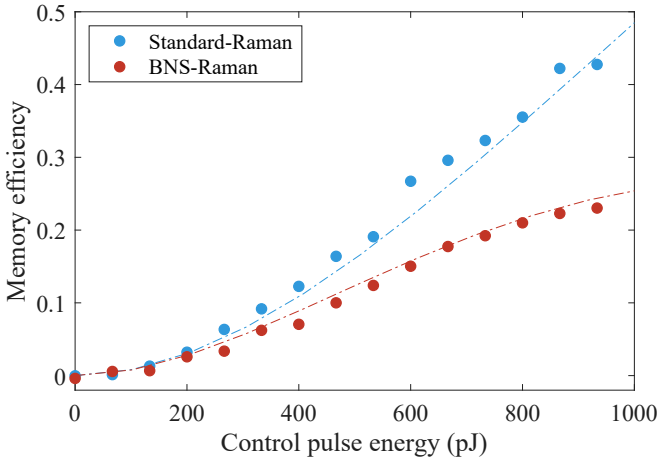


Figure 5.7: Memory efficiency as a function of control pulse energy for the standard (blue) and BNS (red) Raman memory protocols. The memory storage time is 150 ns, and the input photon number for the standard Raman memory is 3.5, and 3.2 for the BNS Raman memory. The dashed lines are the precicted memory efficiency from the numerical simulations and match the data well.

The memory lifetime is (294 ± 26) ns for the standard Raman memory, and (625 ± 36) ns for the BNS Raman memory. The limitation on these timescales is from the extinction of the intensity modulator used to switch the pump light off, causing the spin-wave to be depleted during storage. The discrepancy in the storage time arises from the extinction varying between experiments.

5.3.2 Noise Characterisation

Figure 5.8 shows the average number of noise photons generated per control pulse, N_{noise}, and we can see that the amount of noise is strongly suppressed in the BNS-Raman scheme. At a pulse energy of 670 pJ and storage time of 70 ns the noise level is reduced from $N_{\text{noise}}^{(\text{STD})} = 0.454 \pm 0.002$ to $N_{\text{noise}}^{(\text{BNS})} = 0.0287 \pm 0.0004$ photons per pulse – a decrease by a factor of 16. Note that the errors here are calculated from the Poissonian uncertainty from the number of detection events, and do not account for the uncertainty from, for example, the transmission from the memory to the detectors. The memory efficiency for these parameters is $\eta^{(\text{STD})} = 29.6 \pm 0.1\%$ and $\eta^{(\text{BNS})} = 17.7 \pm 0.1\%$ respectively, and therefore we see a significant decrease from $\mu_1^{\text{STD}} = 1.535 \pm 0.009$ to $\mu_1^{\text{BNS}} = 0.162 \pm 0.002$, where $\mu_1 = N_{\text{noise}}/\eta$. We note that this is very similar to the noise suppression achieved in the first-generation cavity-enhanced Raman memory where $\mu_{1,\text{int}} = 0.17 \pm 0.02$ [207], with a much smaller technical overhead in this present realisation. This demonstrates that the BNS-Raman scheme is a technically simple method to reduce the noise in the Raman memory, without a significant detriment to the memory efficiency or lifetime.

We gain further insight into the built-in noise suppression by looking at how the noise level scales with the memory lifetime. For the standard-Raman memory, the noise in the input time-bin is constant, whilst the noise in the retrieval time bin decreases with storage time with a decay constant of (380 ± 36) ns. This indicates that the noise process involves an atomic coherence that decays over a similar timescale as the memory efficiency – consistent with four-

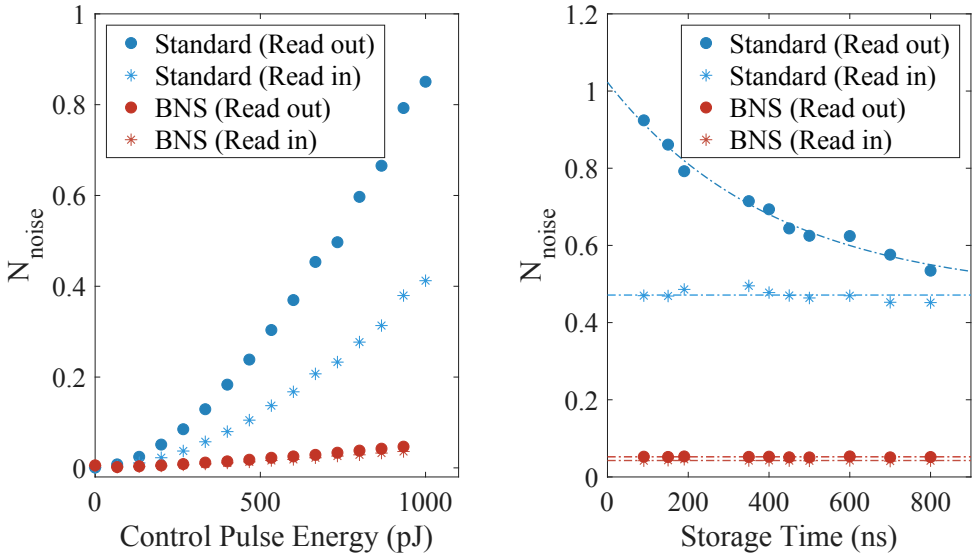


Figure 5.8: The measured noise photons per pulse, N_{noise}, in the input and storage time-bins as a function of control pulse energy (left) and memory storage time (right) for the BNS (red) and standard (blue) Raman memory. The error bars are within the marker size. The left panel was measured using a memory storage time of 70ns, and the right panel was measured with a control pulse energy of 930pJ.

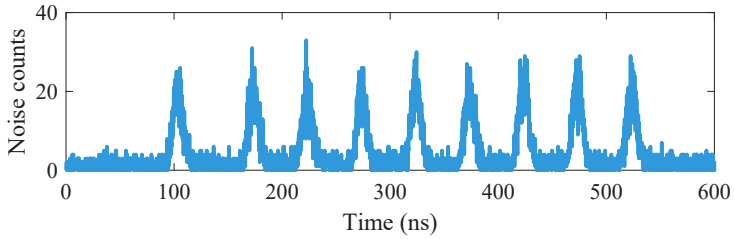


Figure 5.9: The measured noise in the BNS-Raman memory for eight read-out pulses, when the duration of the control pulse is 10 ns.

wave mixing noise. In contrast, for the BNS-Raman memory the noise is independent of the storage time, indicating that we have suppressed four-wave mixing and have some residual technical noise. We also see that in the standard-Raman memory, for the same control energy and storage time the noise is significantly higher in the retrieval time-bin that in the input time-bin. This is again consistent with four-wave mixing noise which builds up with subsequent applications of the control pulse due to the generation and partial retrieval of a spin-wave excitation with each pulse [183]. Other sources of noise such as spontaneous scattering due to imperfect optical pumping and fluorescence noise would generate the same amount of noise photons on every application of the control pulse and therefore would be equal in the storage and retrieval time-bins. We highlight this by measuring the noise for 8 read-out pulses for the BNS-Raman memory, as shown in Figure 5.9, and we see that the noise is almost identical across all time-bins, which again indicates the suppression of four-wave mixing noise.

5.4 Correlation Functions

We investigate the residual noise in more detail by considering the second-order autocorrelation function of the output photonic state. This is defined as:

$$g^{(2)}_{\text{out}} = \frac{\iint \mathrm{d}t\,\mathrm{d}t' \left\langle S^{\dagger}_{\text{out}}(t')S^{\dagger}_{\text{out}}(t)S_{\text{out}}(t)S_{\text{out}}(t')\right\rangle}{\left(\int \mathrm{d}t \left\langle S^{\dagger}_{\text{out}}(t)S_{\text{out}}(t)\right\rangle\right)^2},$$

where S_{out} is the output signal mode operator, and $g^{(2)} < 1$ signifies a field with non-classical statistics. We consider three noise processes that could contribute to the photon number statistics.

- **Spontaneous Raman Scattering**, N_{SRS}, from four-wave mixing noise or spontaneous Stokes scattering from the unpumped thermal population in $|3\rangle$ with $g^{(2)}_{\text{SRS}} = 2$. These noise processes are both contained in the equations of motion, and we can predict their contributions using Green's function analysis.
- **Broadband collision-induced fluorescence**, N_{F}, that is not sufficiently filtered from detection. We note that the noise due to fluorescence will have multimode thermal statistics with $1 \leq g^{(2)}_{\text{F}} \leq 2$, depending on how many modes are collected. The case $g^{(2)}_{\text{F}} = 2$ represents single-mode thermal statistics i.e. collecting fluorescence from a single spatial and temporal-spectral mode, and $g^{(2)}_{\text{F}} = 1$ is fully multimode. We initially assume that the fluorescence noise has single-mode thermal statistics with $g^{(2)}_{\text{F}} = 2$, since our detection is after significant spectral filtering and within a time window of 35 ns which is comparable to the fluorescence timescale, although this is something we can relax later.
- **Control field leakage**, N_{L}, with $g^{(2)}_{L} = 1$. We remove the caesium vapour cell and measure the rate of detection events which is within the dark counts of the detectors, and therefore we deduce that we can neglect this source of noise and set $N_L = 0$.

By taking into account these different noise processes, T.M.H. and O.L.A. derived an expression for the $g^{(2)}$ of the retrieved state as a function of the output photon number $N_{\text{out}} = \eta N_{\text{in}}$, details of which are given in Appendix A. We predict:

$$g^{(2)}_{\text{out}} = 1 + \frac{aN^2_{\text{out}} + 2N_{\text{SRS}}N_{\text{out}} + N^2_{\text{SRS}} + N^2_{\text{F}}}{(N_{\text{out}} + N_{\text{SRS}} + N_{\text{F}})^2} \tag{5.7}$$

where $a = g^{(2)}_{\text{in}}\mathcal{G}_{ss}/\eta^2 - 1$, and $\mathcal{G}_{ss}$ is the integrated Green's function kernel describing the linear mapping from the input signal field to the retrieved signal field. We note that $\mathcal{G}_{ss}$ could be calculated numerically or, as we see later in Chapter 6 measured explicitly, but in this current investigation we keep it as a parameter since our numerical simulations are computationally slow at the BNS-Raman condition and it is experimentally challenging to measure directly.

To quantify the noise processes further we measure the second-order autocorrelation of the retrieved optical field for coherent state inputs with average photon numbers varying from

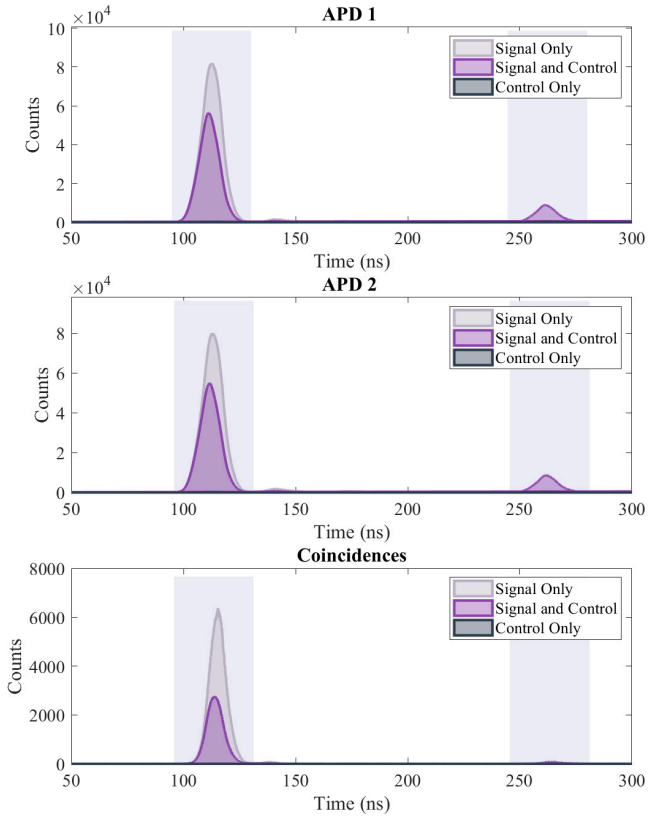


Figure 5.10: Histograms of the photon detection events on the two avalanche photodiode detectors (APD1 and APD2), and the coincidences between them. The grey window shows the coincidence window which is set to be 35 ns.

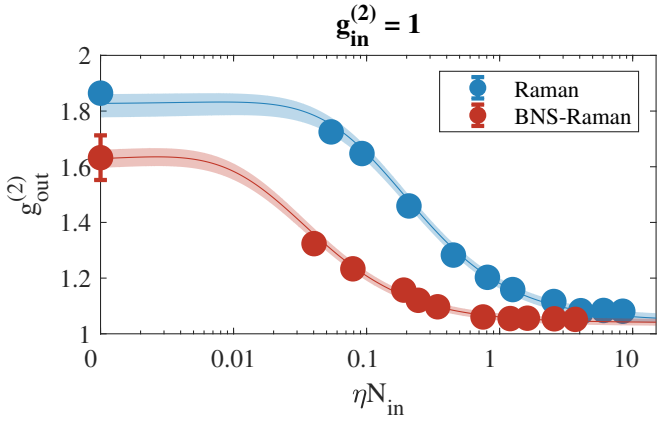


Figure 5.11: The measured $g^{(2)}$ of the retrieved state from the memory as a function of the retrieved photon number, N_{out} for the BNS and standard Raman memory. The control pulse energy is 330 pJ and the storage time is 150 ns. The solid lines are the fit to data using Equation 5.7, and the shaded regions indicate the 95% confidence intervals on the fit.

$N_{\text{in}} \sim 0.5$ to 80. The output signal is sent to a Hanbury-Brown-Twiss set-up comprised of a half waveplate, a polarising beam splitter, and two fibre-coupled APDs. The correlations between the two detectors are measured using a time-tagger (Swabian Instruments Time Tagger 20) to calculate $g^{(2)}_{\text{out}}$. The software used to record coincidences and single detection events in the relevant time windows was written by myself and D.J.S, and a set of example histograms is shown in Figure 5.10. We fixed the coincidence window to be 35 ns to capture all the photons arriving within a pulse which has an intensity full-width half maximum of 10 ns, without counting the photons arriving approximately 30 ns after the pulse due to ringing of the RF amplifiers (see Section 6.2.1). We set the storage time of the memory to be 150 ns for all the autocorrelation measurements to ensure that there is no cross-talk between the input and output pulses when recording coincidences. To detect enough coincidences to ensure a low error on the $g^{(2)}$ measurements we have to acquire data for many hours, particularly when the input photon number is small.

Figure 5.11 shows the results for the measured $g^{(2)}_{\text{out}}$ as a function of the output photon number, $N_{\text{out}} = \eta N_{\text{in}}$, in both memory configurations. BNS-Raman outperforms the standard Raman memory for all input photon numbers tested, with the measured $g^{(2)}_{\text{out}}$ being lower overall and scaling faster to unity as the input photon number is increased. In both cases in Figure 5.11 we see that $g^{(2)}_{\text{out}}$ decreases towards $g^{(2)}_{\text{out}} \approx 1.05$ with increasing photon number, as the relative contribution of the noise photons compared to the retrieved signal decreases and there is less pollution of the photon number statistics. We find that even at high input photon numbers, the retrieved $g^{(2)}_{\text{out}}$ is larger than the input $g^{(2)}_{\text{in}} = 1$ due to some gain in the system which adds thermal statistics to the output and is fully captured within the $\mathcal{G}_{ss}$ Green's function. Since this term is multiplied by $g^{(2)}_{\text{in}}$ in Equation 5.7, we note that this will not contribute to the photon number statistics for a single photon input with $g^{(2)}_{\text{in}} = 0$.

We fit the data according to Equation 5.7. We calculate $\eta = N_{\text{out}}/N_{\text{in}}$ from the memory traces by integrating over the input and output time bins to extract the number of input/output photons and subtracting the noise counts. We also measure the number of detection events in the retrieval time bin when $N_{\text{in}} = 0$, and by measuring the transmission from the memory to the detector we can get an estimate of the total number of noise counts per pulse N_{noise}. We constrain the fit parameters such that $N_{\text{SRS}} + N_{\text{F}} = N_{\text{noise}}$. We therefore have only two free parameters in our fit: $\mathcal{G}_{ss}$, and the relative proportion of noise that is from N_{F} and N_{SRS}. The fit to the data, and the 95% confidence intervals are shown in Figure 5.11, and the extracted fit parameters are given in Table 5.1. We find that N_{SRS} is reduced from $N^{\text{STD}}_{\text{SRS}} = 81(2) \times 10^{-3}$ to $N^{\text{BNS}}_{\text{SRS}} = 11.0(5) \times 10^{-3}$ photons per pulse via the BNS scheme, with N_{F} seeing a slight decrease from $N^{\text{STD}}_{\text{F}} = 9(3) \times 10^{-3}$ to $N^{\text{BNS}}_{\text{F}} = 3.8(5) \times 10^{-3}$. Therefore, along with the significant decrease in the average noise, the BNS case presents a change in the ratio of SRS-to-fluorescence noise (~ 3, compared to ~ 9 for the standard Raman memory). This results in a noticeable reduction of $g^{(2)}_{\text{out}}$ at zero input, and the increased SNR allows for faster scaling to unity as the input photon number is increased.

To distinguish whether the remaining noise from spontaneous Raman scattering, N_{SRS}, is due to four-wave mixing or Stokes scattering due to imperfect optical pumping, we measure the noise as a function of the optical pumping efficiency. The results, shown in Figure 5.12,

	$\mathcal{G}_{ss}(10^{-3})$	$N_{\mathrm{SRS}}(10^{-3})$	$N_{\mathrm{F}}(10^{-3})$
STD-Raman	11.0(2)	81(2)	9(3)
BNS-Raman	10.9(1)	11.0(4)	3.8(5)

Table 5.1: Fit parameters for $g^{(2)}$ of the retrieved state as a function of photon number, according to Equation 5.7. The numbers in brackets give the standard error on the last digit.

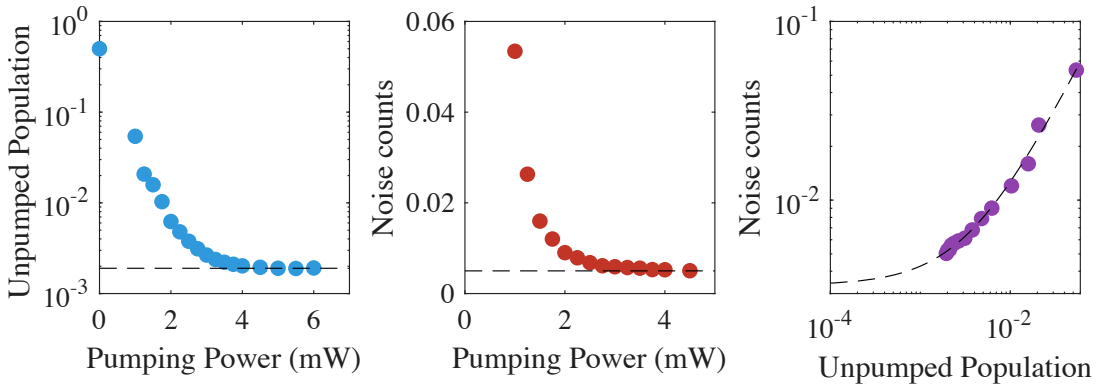


Figure 5.12: (a) The proportion of the population in the storage state, or one minus the optical pumping efficiency, as a function of the power of the optical pumping beam. (b) The number of noise photons per pulse as a function of the pumping power. (c) The number of noise photons per pulse as a function of the amount of unpumped population, i.e. combining graphs (a) and (b).

demonstrate that the noise increases linearly with the amount of residual population in the storage state. However, these data were taken with the optical pumping beam turned on all the time, and not switching off for memory storage. Due to the degradation of the switching EOM we can now only achieve a pumping power of 2 mW before storage, and an extinction ratio of 20:1 during storage[1]. We can extrapolate the data in Figure 5.12 to predict that, even if we had perfect optical pumping, we would still have a noise level of $N_{\mathrm{noise}} = (5 \pm 2) \times 10^{-3}$. This is very similar to the extracted $N_{\mathrm{F}}^{\mathrm{BNS}} = 3.8(5) \times 10^{-3}$ in Table 5.1, and we therefore postulate that the majority of the remaining noise is due to imperfect optical pumping and fluorescence noise, and that we have successfully eliminated four-wave mixing noise. We can see in Figure 5.12 that we could reduce the noise from $N_{\mathrm{noise}} = 11.8(2) \times 10^{-3}$ to $N_{\mathrm{noise}} = 6.6(2) \times 10^{-3}$ by increasing the optical pumping power from 2 mW (as was the case for the autocorrelation data in Figure 5.11) to 5 mW, and potentially even lower if we can achieve a pumping efficiency greater than the maximum 99.8% demonstrated in Figure 5.12. We intend to replace the EOM with a Pockels cell which would allow for high transmission and high extinction of the optical pumping beam to enable a reduction in the residual noise.

If we can successfully reduce the noise due to imperfect optical pumping, then the fluorescence noise, N_{F}, will become the next biggest noise source. There are several ways we can consider reducing the noise due to fluorescence. Firstly, we note that our filtering set-up is not optimal for this new operational detuning of the memory. The first four Fabry-Pérot etalons that we use to filter the control field and other sources of noise have a free-spectral range of 18.4 GHz. This was chosen to maximally suppress the control field which is 9.2 GHz detuned

[1]This is in contrast to an extinction ratio of 200:1 that we mentioned in Chapter 3. The EOM has suffered photorefractive damage and is no longer performing nearly as well as in previous experiments

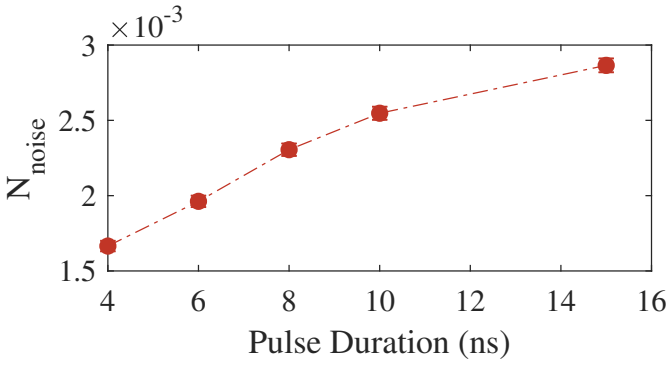


Figure 5.13: Noise photons per pulse, N_{noise}, as a function of the full-width at half maximum pulse intensity duration. The pulse energy was kept fixed at 50 pJ and the memory efficiency was $(0.5 \pm 0.1)\%$ in all measurements.

from the signal. However, when the memory detuning is $\Delta = -2\Delta_{\text{hf}} = -18.4$ GHz, these filters have a transmission peak on resonance with the atomic transition. The majority of fluorescence noise will be emitted close to resonance and therefore will have a high transmission through these filters. The fluorescence noise is only being significantly filtered by the final two etalons with an FSR of 103 GHz. We intend to change the filtering set-up to use etalons with an FSR of 40 GHz instead of 18.4 GHz, which we hope will significantly reduce the fluorescence noise.

We could also reduce the amount of fluorescence noise in the same time window as the retrieved pulse by decreasing the pulse duration. We discussed earlier in Chapter 4 that fluorescence noise is emitted over a characteristic time-scale of the excited state which is around 30 ns, meaning that the majority of the fluorescence noise is captured within our 35 ns integration window. We measured the amount of noise photons per pulse as a function of pulse duration, whilst keeping the mean pulse energy fixed. The results in Figure 5.13 show that the noise decreases for shorter pulse durations, whilst the memory efficiency was constant at $(0.5 \pm 0.1)\%$ for these data. The efficiency is so low because the pulse energy is confined by the shortest pulse duration, and is therefore fixed at 50 pJ. By carving shorter pulses with high pulse energy we would be able to decrease the fluorescence noise that is captured within the pulse duration.

The final way to reduce fluorescence noise would be to decrease the amount of linear absorption of the control field, which leads to fluorescence noise. With the current operational temperature and buffer gas pressure, the control field experiences around 16% linear absorption at the BNS detuning condition. If we removed the buffer gas entirely this would decrease to 1%, and therefore we would expect significantly lower fluorescence. But it is not as straightforward as simply removing the buffer gas, since we saw in Chapter 3 that the buffer gas plays a vital role in achieving a high optical pumping efficiency, and therefore there is a compromise between N_{F} and N_{SRS}. However, it may be possible to reach a parameter regime where the total amount of noise due to imperfect optical pumping and fluorescence is minimised.

We can now use Equation 5.7 together with the fitting parameters from the weak coherent state data to predict the output photon statistics for the case when the input is a single photon Fock state with $g^{(2)}_{\text{in}} = 0$. Figure 5.14 shows the predicted $g^{(2)}_{\text{out}}$ as a function of probability of heralding a single photon at the input of the quantum memory η_{h}. The standard Raman memory is unable to produce a non-classical output state, even for unit heralding, due to the

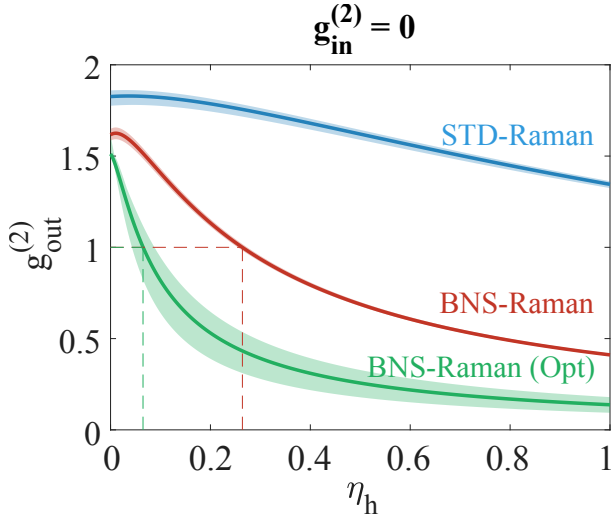


Figure 5.14: The predicted $g^{(2)}_{\text{out}}$ for a single photon input with $g^{(2)}_{\text{in}} = 0$ as a function of the input photon number or heralding efficiency, η_h, using the fitted parameters. The line in green is a prediction for the output $g^{(2)}$ if we could decrease noise due to spontaneous Raman scattering by a factor of two and remove fluorescence noise.

significant noise contribution. In contrast, the BNS-Raman case with the large reduction FWM noise is able to output non-classical states for heralding efficiencies exceeding $(26.4 \pm 0.5)\,\%$ – representing a significant improvement by simply changing the detuning of the memory protocol.

If we implemented the changes outlined above, we would be able to improve further on these results. By increasing the optical pumping power we predict that we can decrease N_{SRS} by almost a factor of two, and by improving the filtering we expect to be able to significantly reduce the fluorescence noise. We could also increase the lifetime of the memory by improving the low switching extinction on the pumping mode, back to the μs timescales demonstrated in Chapter 3, and this would increase the memory efficiency at a storage time of 150 ns from 10.2 % to 12.7 % without increasing the noise. These changes would give $\mu_1^{\text{opt}} = 0.052(3)$, or an upper bound on the conditional fidelity of a retrieved qubit of around $\mathcal{F} = 0.95$ [157]. This would also yield the green line in Figure 5.11(b), or a drop in the requisite single-photon heralding efficiency for non-classical read-out to $(6.5 \pm 0.6)\,\%$. In this case a photon from a deterministic single photon source could be retrieved from the memory with $g^{(2)}_{\text{out}} = 0.14 \pm 0.02$, significantly below the non-classical threshold. Further improvements could be made by shaping the control field to optimise the memory efficiency for a given pulse energy, thereby not introducing any further noise. We hope that with these further improvements, single photons will be able to be retrieved with very low $g^{(2)}$ in the near future using this BNS scheme.

5.5 Summary

In this chapter I have presented a new scheme for the suppression of four-wave mixing noise using atomic absorption, and the key results are outlined below.

- In Section 5.1 I introduced the built-in noise suppression (BNS) scheme for suppressing four-wave mixing noise. This scheme uses atomic absorption to suppress unwanted scat-

tering, and is broadly applicable to any off-resonant memory protocol such as Raman, DLCZ, and $\Lambda-$GEM.

- In Section 5.2 I presented our numerical simulations of the BNS scheme, which demonstrate over four orders of magnitude suppression in four-wave mixing type noise.
- In Section 5.3 I investigated the BNS scheme experimentally, and observed an order of magnitude reduction in the noise performance. The noise-to-signal ratio quantified by $\mu_1 = N_{\text{noise}}/\eta$ decreased from $\mu_1^{\text{STD}} = 1.535 \pm 0.009$ to $\mu_1^{\text{BNS}} = 0.162 \pm 0.002$.
- In Section 5.4 I characterised the residual noise by measuring the second-order autocorrelation of the retrieved state from the memory, $g_{\text{out}}^{(2)}$. I measured $g_{\text{out}}^{(2)}$ as a function of the output photon number, and fitted the data according to a model developed by T.M.H. and O.L.A. The fit parameters suggest a decrease in noise due to spontaneous Raman scattering from $81(2) \times 10^{-3}$ to $11.0(5) \times 10^{-3}$. I used this fit to predict non-classical recall when storing heralded single photons with a heralding efficiency of more than $(26.4 \pm 0.5)\,\%$. We predict that the majority of the residual noise is due to imperfect optical pumping which we could improve by increasing the power of the pumping beam. We could also reduce the fluorescence noise by improving the filtering set-up. This would enable non-classical recall from the memory with a heralding efficiency greater than $(6.5 \pm 0.6)\,\%$.

This scheme is widely applicable to any system that suffers from four-wave mixing noise, and in particular this method would be very efficacious in a cold atomic system where collisional induced fluorescence noise is negligible. This is a technically simple and effective noise suppression scheme, that paves the way towards quantum level storage in a long-lived, broadband, room-temperature quantum memory.

Chapter 6:

Temporal Mode Properties of the Raman Memory

In this final chapter we investigate a different property of the Raman memory, namely its modal properties. We discussed in Chapter 2 that the Raman memory can be approximated as single mode in the low coupling regime, and in Chapter 4 we mentioned that the cavity enhances the single-modeness. Here we investigate this experimentally for the first time. I will first introduce temporal modes and their applications in more detail, and then demonstrate how the Raman memory can be used as a device to separate and manipulate temporal modes of light.

This project was overseen by B.B. and the pulse carving infrastructure and software was developed by J.H.D.M. I built the new memory set-up for interfacing the temporal modes with the Raman memory, including the pulse amplifier, filtering and detection. J.H.D.M. and I contributed equally to the tomography experiment, whilst the theoretical analysis was performed by J.H.D.M. and J.S. I conducted the temporal mode manipulation experiments in the final section of this chapter.

6.1 Temporal Modes

Temporal modes (TMs) are field-orthogonal broadband wavepacket states of pulsed light which span an infinite-dimensional Hilbert space. In this chapter we explore two main motivations for studying TMs of quantum light: firstly to enable high-dimensional encodings in photonic quantum information processing, and secondly to enhance light-matter interactions by optimally shaping the temporal-spectral mode of light. To harness the potential of TMs in this way we need to be able to address a single mode and also convert between different TMs. The Raman memory offers this capability: the single-mode nature of the memory allows us to address only a single TM, and by shaping the control pulse that drives the interaction we can determine which mode is addressed. Additionally, by selecting distinct modes for the read-in and read-out

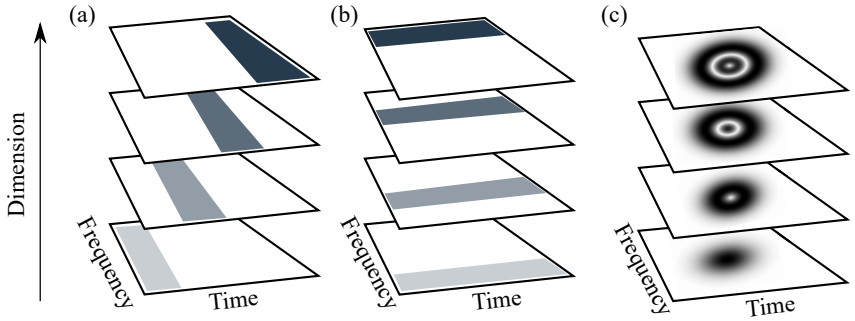


Figure 6.1: High dimensional temporal-mode encodings in time-frequency space. Orthonormal bases can be constructed using time bins (a) or frequency bins (b), or by using field-orthogonal pulsed temporal modes such as Hermite Gaussian modes (c). Figure adapted from [219].

control pulses we can convert the input signal into a different TM. This novel platform for TM operation offers the means to separate, delay, and manipulate user-defined temporal modes in the MHz–GHz regime in a single device – a key capability for future quantum technologies.

6.1.1 Temporal Modes as a Basis for QIP

Temporal modes have recently been highlighted as an appealing basis for quantum information processing [214]. We discussed in Section 1.2.6 that photons are a promising platform for quantum communication as they are noise-free at room temperature and have many degrees of freedom in which to encode information. The majority of work on photonic information processing has thus far focused on polarisation encodings, as it is easy to prepare, manipulate, separate and measure the photonic state using polarisation optics. However, polarisation is fundamentally limited to a two-dimensional Hilbert space where information is encoded in qubits, and only one bit of information can be transmitted per photon using the polarisation basis alone. There has been much recent interest in using higher dimensional bases for photonic encoding to increase the information capacity of communication links, and such *qudit* bases include transverse spatial modes and time-frequency bins [215, 216].

Encoding information in spatial qudits such as orbital angular momentum (OAM) states has enabled demonstrations of high dimensional quantum key distribution protocols [217, 218]. However, multiple OAM modes are not compatible with single-mode fibres and have to be transmitted via free space links, which limits the distance over which they can be transmitted. OAM states also suffer from perturbations such as dispersion and air turbulence.

There is growing interest in using time-frequency modes as photonic qudits, since they overlap spatially and are hence compatible with single-mode optical fibres and waveguide devices, and all modes are identically affected by dispersion and spatial perturbations. A simple time-frequency encoding uses discrete time or frequency bins which form a quasi-orthogonal basis as shown in Figure 6.1(a)-(b). However, these modes are not optimally packed in time-frequency space. We therefore consider temporal modes, which overlap in time and frequency, but are field-orthogonal and form a complete orthonormal basis to encode arbitrary states, as shown in Figure 6.1(c). TMs achieve optimal packing in time-frequency space and hence maximise the achievable dimensionality for a given detector bandwidth, in an analogous way to Code-division Multiple Access (CDMA) encoding at RF frequencies for telecommunications [220].

However, since TMs overlap spatially, temporally and spectrally, it is challenging to address individual basis states and perform qudit rotations or measurements in this basis. In analogy to the polarisation basis where we can use polarising beam splitters to separate basis states, and waveplates to manipulate them, we instead need a time-dependent beam splitter that can separate and manipulate the temporal waveform of light [221].

6.1.2 Enhancing Light-Matter Interactions

In addition to using temporal modes as a high dimensional basis for QIP, we can consider tailoring the temporal mode of light to enhance the interaction between light and material systems. As discussed in Chapter 1, future quantum technologies will harness the advantages of multiple quantum systems in a hybrid quantum network – combining high fidelity information processing at matter-based computational nodes connected via fast, long-distance photonic links. However, atomic systems typically have acceptance bandwidths on the order of 10s–100s MHz, while photonic communication links typically operate at 10s of GHz, and to optimally interface these systems we require the capability to efficiently and coherently convert the bandwidth of light in the MHz and GHz regime. Furthermore, an atomic two-level system will emit light with a falling exponential wavepacket, and conversely will optimally absorb light with a rising exponential wavepacket [222]. To interface two atomic systems it is advantageous to be able to interconvert between rising and falling exponentials, amongst other temporal modes of light.

In addition, we consider how controlling the temporal mode of light can directly enhance light-matter interactions in the context of the Raman memory. In Chapter 3 we discussed how the Raman interaction strength increases with the Rabi frequency of a strong control field, but the AC Stark shift destroys the two-photon resonance condition and the memory efficiency is limited. By dynamically tailoring the amplitude and phase of the control field we can compensate for the AC stark shift and maintain the two-photon resonance condition, throughout the memory interaction, enabling significantly higher memory efficiencies. Optimising the memory efficiency by shaping the control pulse has been demonstrated in an EIT memory [223] and, more recently, in a Raman memory [179].

The ability to coherently manipulate TMs, for example by adjusting the bandwidth or complex wavepacket, is therefore a key capability of future quantum technologies.

6.1.3 Devices for Manipulating Temporal Modes

To use temporal modes as a high-dimensional basis for photonic information encoding, and to enhance light-matter interactions by tailoring the wavepacket of light, we need to be able to prepare, manipulate and separate TMs. To operate on overlapping, field-orthogonal temporal modes we require a time non-stationary beam splitter interaction. In general such interactions involve a light-matter interface which is driven by a time-varying external control field. There are two main candidates for TM devices: the quantum pulse gate based on nonlinear optical interactions [71], and a single mode quantum memory [190, 205].

6.1.3(a) Quantum Pulse Gate

The quantum pulse gate (QPG), first proposed in 2010 [224], uses sum frequency generation (SFG) to manipulate temporal modes. A signal and control beam are passed through a non-linear medium, and part of the signal field is frequency converted in a $\chi^{(2)}$ process. The control field profile determines which temporal mode of the signal is addressed, and the selected TM is converted into a different wavelength and can subsequently be separated with standard optical components. The QPG requires careful dispersion engineering such that the group velocity of the signal and control fields are matched. A quantum pulse gate in lithium niobate waveguides has been realised [71] and recent results put forward ways of addressing single TMs with high fidelities [225]. However, such TM-manipulation devices are limited to few-picosecond or shorter pulses (bandwidths $\gtrsim$ 10s GHz) due to the stringent phase-matching conditions [226], which is not compatible with many light-matter interfaces which typically operate in the MHz–GHz regime.

6.1.3(b) Single Mode Quantum Memory

A single-mode quantum memory can enable TM manipulation in the MHz–GHz bandwidth regime. As described in Chapter 2 the ideal Raman memory interaction is a time-dependent beam splitter that links optical temporal modes to matter spin-waves, and hence is a new candidate for temporal mode operation. The ideal (noise-free) Raman memory is described by a set of linear equations between the signal mode S, and the spin-wave mode B, driven by read-in and read-out control pulses $\Omega_{\text{in/out}}(t)$. We discussed in Section 2.2.2 that, since these equations are linear, the resulting evolution may be characterised by a Green's function mapping the input signal, S_{in}, to the storage mode, B_{stor}:

$$B_{\text{stor}}(z) = \int_{-\infty}^{+\infty} \mathrm{d}t'\, K_1(z,t') S_{\text{in}}(t'), \tag{6.1}$$

where K_1 encodes the details of the read-in control pulse and fully describes the memory storage interaction. We can similarly describe the retrieval process of the memory as:

$$S_{\text{out}}(t) = \int_{0}^{1} \mathrm{d}z'\, K_2(z',t) B_{\text{stor}}(z'). \tag{6.2}$$

It can be shown that in the low coupling regime, the Raman memory is approximately single mode and the storage Green's function K_1 maps a single input temporal mode to a stored spin-wave [205]. Any mode orthogonal to this principal eigenmode will be unaffected, and transmitted through the memory. Similarly the retrieval interaction can also be approximated as single mode where the spin-wave will be mapped (via K_2) to a single output temporal mode which is determined by the retrieval control pulse $\Omega_{\text{out}}(t)$.

The Raman memory presents an ideal candidate for TM manipulation since it operates in the MHz–GHz bandwidth regime, and the single-mode nature can be used to address a single mode and separate temporally-overlapping TM basis states. Furthermore, by selecting different modes for the read-in and read-out memory interactions we can convert the TM of the input

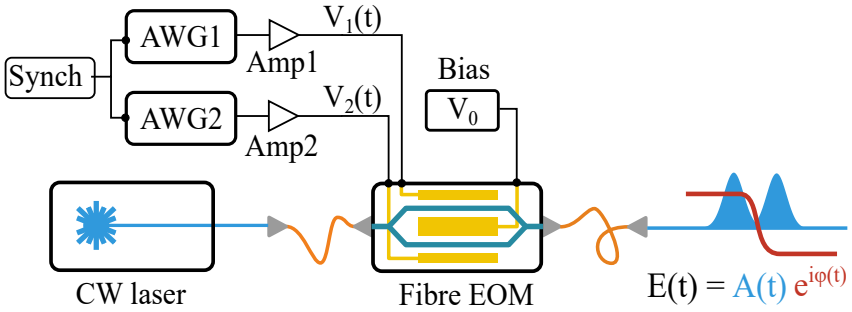


Figure 6.2: A schematic for the pulse carver based on a fibre-integrated Mach-Zehnder interferometer (FMZI), enabling arbitrary amplitude and phase control. Two arbitrary waveform generators (AWGs) are synchronised and used to generate time-varying RF signals, which are then amplified using RF amplifiers. These signals, $V_1(t)$ and $V_2(t)$ are applied to each arm of the interferometer to modulate the phase accrued by the light. A DC bias voltage is applied between the two waveguides in the interferometer. The amplitude of the output pulse, $A(t)$, is governed by $V_1(t) + V_2(t)$, and the phase, $\phi(t)$, is determined by $V_1(t) - V_2(t)$.

optical signal into a user-defined output mode.

In this chapter we investigate the modal properties of the Raman memory and its potential application as a device for manipulating temporal modes. We first quantify the degree to which the Raman memory is mode selective, by tomographically reconstructing the memory interaction and quantifying the number of eigenmodes of the system. Once we establish that the Raman memory is single mode we then demonstrate conversion between different temporal modes.

6.2 Temporal Mode Raman Memory Set-up

To investigate the temporal mode properties of the Raman memory we built a new set-up where we can independently control the temporal mode of the signal and control pulses. The pulses are generated using a fibre-integrated Mach-Zehnder interferometer to carve pulses from a continuous wave laser, rather than from a pulsed Ti:Sapph laser (as was the case for the previous Raman memory outlined in Chapter 3). The pulses from the pulse carver have to be amplified to have sufficient energy to drive the memory interaction. The pulse carver infrastructure and software interface was built by J.H.D.M, and the amplifier and the rest of the Raman memory set-up were built by myself.

6.2.1 Pulse Carver

To carve arbitrary pulses of light we use a fibre-integrated Mach-Zehnder interferometer (FMZI) where we apply a time-varying phase to each path, enabling full control of the amplitude and phase of the transmitted light. A schematic of the FMZI is shown in Figure 6.2. Applying voltages $V_1(t)$ and $V_2(t)$ to the two arms of the FMZI electro-optically changes the refractive index of the nonlinear medium and the light undergoes a phase shift which is linearly proportional to the applied voltage. The light is subject to phases $\phi_1(t) = \alpha(V_1(t) - V_0)$ and $\phi_2(t) = \alpha(V_0 - V_2(t))$, where α is determined by the response of the device, and V_0 is a constant

bias voltage applied between the two arms of the interferometer. The resultant field is:

$$\begin{aligned} E_{\text{out}}(t) &\propto \cos\left(\alpha\frac{V_1(t)+V_2(t)}{2}+\alpha V_0\right)\exp\left(\mathrm{i}\alpha\frac{V_1(t)-V_2(t)}{2}\right)\exp(\mathrm{i}\omega_0 t) \\ &= A(t)e^{\mathrm{i}(\omega_0 t+\phi(t))}. \end{aligned} \tag{6.3}$$

We can set V_0 such that $A(V_1 = V_2 = 0) = 0$ and $A(t) = \sin\left(\alpha\frac{V_1+V_2}{2}\right)$. The amplitude of the output field, $A(t)$, is modulated by $V_1(t) + V_2(t)$, and the phase of the pulse, $\phi(t)$, is controlled by $V_1(t) - V_2(t)$, so we can create any arbitrary pulse of light.

To carve pulses we use two arbitrary waveform generators (AWGs) with a sampling rate of 50 Gs/s to generate the necessary time-dependent voltage signals. The AWGs (Tekronix AWG70000) are synchronised (Tektronix AWGSynch) and amplified by two 10.7 Gb/s Driver Amplifiers (PicoSecond Pulse Labs 5868) to achieve sufficient amplitude to drive the FMZI. We have independent control of the gain setting and zero-point offset of the two RF amplifiers, and therefore we can ensure that the two voltage signals are identical in amplitude. The bandwidth of the pulses we can generate is limited by the bandwidth of the RF amplifiers, and we see significant ringing of the amplifiers after the pulse if we drive them too quickly, and this becomes noticeable if we attempt to carve pulses shorter than a nanosecond in duration. This system is nonetheless sufficient to generate ns-duration pulses with non-trivial amplitude and phase control.

The DC bias voltage, V_0 can be varied between 0 and V_π to adjust the overall transmission of the pulse carver, where $V_\pi = \pi/\alpha$ quantifies the response of the EOM. We found that this voltage had to be continually monitored and adjusted due to thermal effects and motion of charge carriers within the EOM changing the V_π value of the device. We therefore applied a time-varying signal that alternates between $+V_\pi$ and $-V_\pi$ at a frequency of 10 Hz to minimise the macroscopic motion of the charge carriers and keep V_π from changing; a technique known as "symmetric biasing". We also use a photodiode to monitor a small fraction of the output light when $V_1 = V_2 = 0$ and actively feed back on the bias voltage to minimise this value and stabilise the system.

6.2.2 Tapered Amplifier

The FMZI has a maximum CW input power of 40 mW, and therefore to have sufficient pulse energy to drive the Raman memory we use a tapered amplifier (Sacher Serval TEC-400-850-2000) to amplify the pulses after the interferometer. The gain in power $\left(\frac{P_{\text{out}}}{P_{\text{in}}}\right)$ of the tapered amplifier (TA) for a CW input of 11 mW is shown in Figure 6.3(a), and we see that we can increase the power by a factor of almost 200. However, the TA adds noise to the signal via amplified spontaneous emission (ASE), which needs to be filtered so as not to introduce noise to the Raman memory. Furthermore, if the amplifier is operated at currents higher than 1000 mA, corresponding to a gain greater than 25, the output pulse shape can become distorted.

We measure the output of the TA on an optical spectrum analyser (OSA, Anritsu MS9740A), and see that the amplified spontaneous emission is very broadband and spans the whole gain medium from approximately 830 nm to 875 nm, as shown in Figure 6.3(b). The vast majority

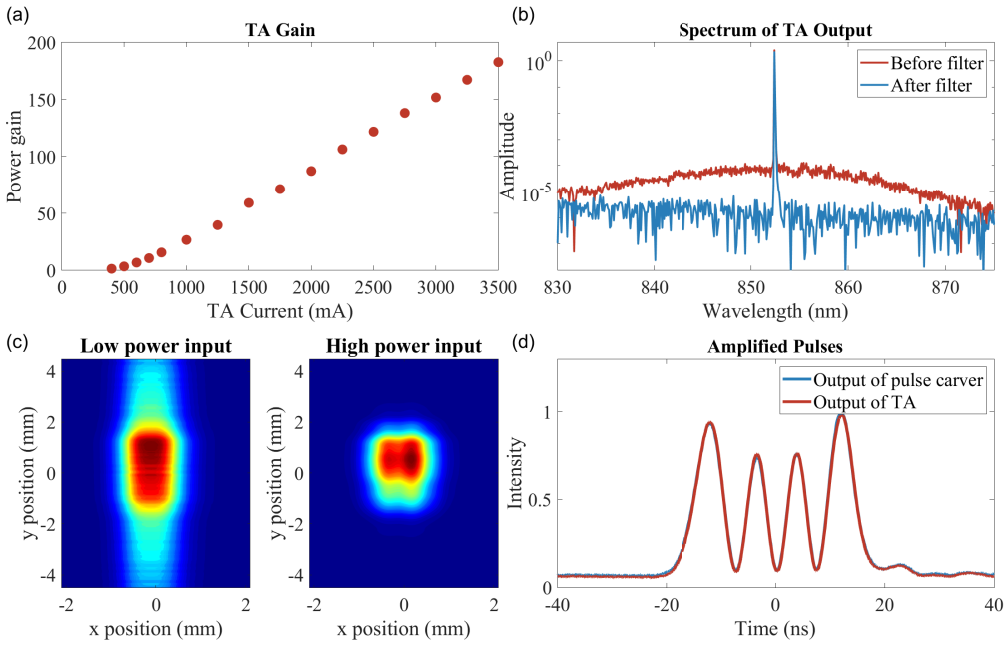


Figure 6.3: Characterisation of the Tapered Amplifier. (a) Gain in power for a CW input laser of 11 mW as a function of TA Current. (b) Output spectrum of the TA measured on a optical spectrum analyser, with and without a holographic grating filtering the amplified spontaneous emission. (c) The output mode of the TA measured on a beam profiler (Thorlabs Dual Scanning Slit Beam Profiler BP209-VIS/M), for a low power (200 μW) and high power (5 mW) input seed. (d) Optical pulses measured before and after amplification, with TA current = 700 mA.

of this noise can be filtered away using a holographic grating. We use a holographic grating which has a measured full-width at half maximum of approximately 60 GHz. The grating has an efficiency of 87% and suppresses all of the broadband noise more than 60 GHz from the seed frequency by several orders of magnitude.

We need to filter the remaining ASE that is within 60 GHz of the control frequency, without losing too much pulse energy. It is particularly important to remove the noise at the signal frequency as this transmits through all of the post-memory filtering (described in Section 3.1.3) and can pollute the output mode of the memory. We use an etalon with a free-spectral range of 103 GHz to isolate the control frequency, where in order to have high transmission through the etalon the beam must be in a collimated, well-defined spatial mode to ensure high interference within the Fabry-Pérot cavity. The output spatial mode directly from the TA is poor and, in particular, when the seed power is low the mode is elongated vertically due to the shape of the gain medium inside the amplifier (see Figure 6.3(c)). We therefore fibre-couple the output of the TA using a single-mode fibre before passing through the etalon, and achieve a transmission of the etalon of 64%. The entire filtering set-up, including the optical fibre coupling of 60%, has a transmission of 33%.

We measure the optical pulses before and after the TA on a fibre-coupled fast photodiode (Thorlabs PDA-8GS) and fast oscilloscope (Teledyne Lecroy Wavesurfer 10). We see in Figure 6.3(d) that in the low-gain regime of the TA, the pulse shape is virtually unchanged by the amplifier. However, as the current of the TA increases the pulses become distorted, as shown in Figure 6.4. There is a compromise between the pulse energy we can achieve and the overlap

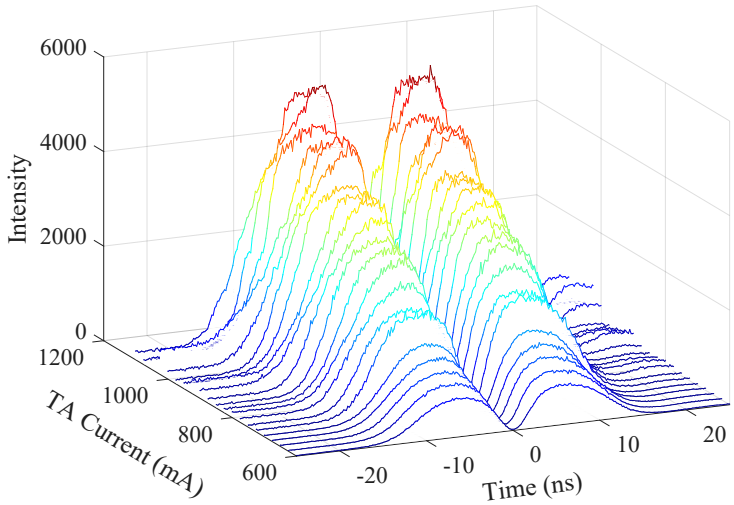


Figure 6.4: Optical pulses measured on APDs as a function of the TA amplification current. At low currents (below 800 mA) the shape of the pulses has good overlap with the ideal second order Hermite-Gaussian, but as the TA current increase the pulses broaden and become distorted.

of the pulses with the ideal temporal mode we are trying to carve. With an input power of 40 mW to the pulse carver, and a gain setting of the TA of around 800 mA we can achieve a pulse energy of several hundred pJ for 10 ns duration pulses which is sufficient to drive the memory interaction with minimal distortion of the pulse shape.

6.2.3 Temporal Mode Raman Memory

A schematic for the experimental set-up is given in Figure 6.5(a). The signal and control pulses are generated using the pulse carving system described above, but the rest of the memory set-up including the optical pumping and post-memory filtering of the control field is the same as that outlined in Chapter 3.

We use the pulse carving set-up to generate three pulses: the signal, read-in control and read-out control, of which we have independent amplitude and phase control. The entire pulse sequence has a repetition rate of 500 kHz, and a schematic of the timing is shown in Figure 6.5(b). The pulses are amplified by the TA, and a small fraction of the amplified pulses is split off and passed through an EOM driven at 9.2 GHz to generate a red sideband at the signal frequency as in Section 3.1.2. The delayed signal is spatially recombined with the control field and the separation between the first two pulses (T_{fib}) is chosen to match the fibre delay such that the attenuated signal and read-in control are temporally overlapped, and the signal is stored in the memory. The separation between the second and third pulses (control read-in and read-out) can be changed to adjust the storage time of the memory.

The pulses are overlapped spatially and focused to a waist radius of 130 μm in the centre of a caesium vapour cell. The vapour is heated to a temperature of 83.0°C to give an optical depth of 3.65×10^4, and placed inside a μ-metal magnetic shield to reduce magnetic dephasing of the spin-wave. A buffer gas of 5 Torr of N_2 is mixed with the caesium vapour to increase the lifetime and allow a high pumping efficiency of $(1 - \alpha) = 99.85\%$ to be reached at such high optical depths, as described in Chapter 3. The filtering of the control field after the memory is the same as in Section 3.1.3.

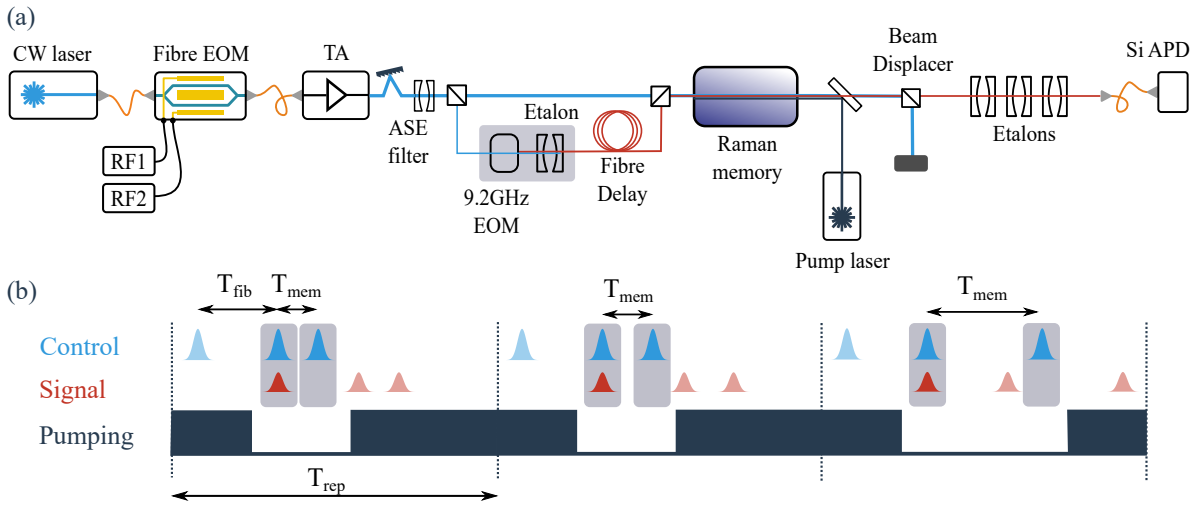


Figure 6.5: (a) A schematic of the temporal-mode Raman memory experiment. Pulses are carved from a CW laser using a fibre-integrated electro-optic interferometer, and amplified using a tapered amplifier (TA). The amplified spontaneous emission (ASE) from the TA is filtered using a holographic grating and a Fabry-Pérot etalon. A small portion of the light is split off and passed through an EOM driven at 9.2 GHz to generate a sideband at the signal frequency, and the carrier frequency and other sideband are filtered away using an etalon. The signal is coupled into a long delay fibre, and then the signal and control beams are overlapped temporally and spatially inside a warm caesium vapour cell. The control field is filtered away using a series of polarisation and frequency filters. The caesium ensemble is prepared via a counter-propagating optical pumping beam. (b) A timing schematic of the signal, control and pumping light. Three pulses are generated by the pulse carver, and the signal mode is delayed by a long fibre delay. The first pulse of the signal is read-in to the memory by the second pulse in the control mode. The final pulse reads in the control mode the signal out from the memory after a user-chosen storage time. The optical pumping is switched off just before storage and remains off throughout the memory interaction.

This set-up allows us to operate the Raman memory with different temporal modes for the signal, read-in, and read-out control pulses, and provides a platform for investigating the temporal mode properties of the Raman memory. In Section 6.3 we first investigate the single mode nature of the memory, and quantify the degree to which the memory can address a single temporal mode. Next, in Section 6.4 we investigate converting between different TMs using this platform.

6.3 Mode Selectivity

We discussed in Section 6.1.3(b) that the Raman memory is single mode in the low coupling regime and will address only a single temporal mode. The mode that is stored by the memory is determined by the read-in control pulse, and any orthogonal modes to this eigenmode will not be stored by the memory and transmit straight through. For an initial, qualitative investigation of the single mode properties of the Raman memory we measure the memory performance when the signal and control have strongly overlapping TMs, and when they have near-orthogonal TMs. The results are shown in Figure 6.6, and we see that when the signal and control TMs are matched, as is the case for the on-diagonal plots, there is significant read-in to the memory since the signal has good overlap with the principal eigenmode of the storage interaction. When

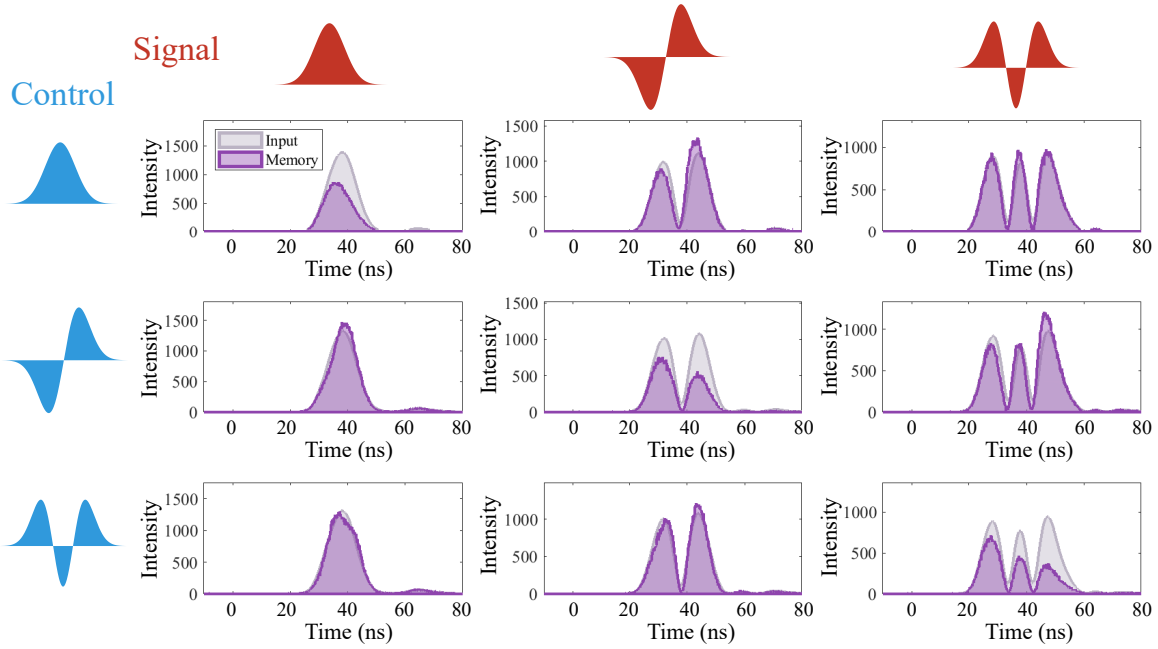


Figure 6.6: A demonstration of mode-selectivity with a single mode Raman memory. The signal is stored in the memory when the signal and control pulses are in the same temporal mode (on-diagonal plots). When the signal and control pulses are orthogonal there is negligible storage in the memory and the input (grey) and memory (purple) traces are virtually identical.

the signal is orthogonal to the TM of the control pulse, the storage efficiency is almost zero and the signal transmits straight through the memory, as shown in the off-diagonal plots.

To quantify the degree of mode selectivity in the Raman memory, we perform a process tomography on the storage interaction. By measuring the read-in efficiency for a tomographically complete set of signal and control modes, we can reconstruct the eigenmodes of the system via a maximum likelihood estimation. This tomography experiment and analysis was led by J.H.D.M, with assistance from B.B. and J.S. I built the memory set-up described above in Section 6.2, and conducted the tomography measurements with J.H.D.M.

As discussed in Section 6.1.3(b), the memory storage interaction is described by a linear mapping from the input signal, S_{in} to a stored spin wave mode B, mediated by a control field $\Omega_{\text{in}}(t)$. We can write this as:

$$B = \mathbf{\Lambda}(S_{\text{in}} \otimes \Omega^*), \tag{6.4}$$

where Λ describes the linear map and acts on the space of signal (S) and control (Ω) modes. We define the storage efficiency in terms of the number of excitations as:

$$\eta_{\text{store}} = \frac{N_B}{N_{\text{in}}} = \frac{\int_0^L \mathrm{d}z\, B^\dagger(z)B(z)}{\int_{-\infty}^{+\infty} \mathrm{d}t\, S_{\text{in}}^\dagger(t)S_{\text{in}}(t)}.$$

$$= (S \otimes \Omega^*)^\dagger \mathbf{\Lambda}^\dagger \mathbf{\Lambda}(S \otimes \Omega^*) \tag{6.5}$$

where we have assumed that the input signal modes are normalised such that the denominator equals one. Therefore, measuring the storage efficiency of the memory allows us to probe the underlying process matrix $\boldsymbol{P} = \boldsymbol{\Lambda}^\dagger \boldsymbol{\Lambda}$. By measuring the storage efficiency for a complete set of optical states S and Ω we can reconstruct a self-consistent process matrix that describes this mapping via a maximum likelihood estimation, and extract information about the modal properties of the Raman memory. In particular, if if $\boldsymbol{P}$ has a single non-zero eigenvalue then the memory storage interaction is truly single-mode and only one particular temporal mode signal will be stored in the memory and all orthogonal modes will be transmitted.

We note that this process tomography method is much more general than the measurement tomography that is usually performed on similar systems, for example the QPG [227]. Typically an interaction is probed using a small set of control modes, for example the first few Hermite-Gaussian modes, and a complete set of signal modes. This tells you about how single mode the operation is *for these control modes*. Our method allows us to go beyond this, and investigate how single mode the operation is *regardless of the control mode*. Furthermore, once the process matrix $\boldsymbol{P}$ is known, we can select the particular control mode, Ω, such that the projection of $\boldsymbol{P}$ onto this mode, $\boldsymbol{P}_{\boldsymbol{\Omega}} = (\mathbb{1} \otimes \Omega^*)^\dagger \boldsymbol{P} (\mathbb{1} \otimes \Omega^*)$ is any chosen signal mode: by tailoring the control field, we can choose what the principal eigenmode of the memory interaction is, such that it matches our chosen input signal mode.

We perform a process tomography of the storage interaction of the Raman memory in the three-dimensional Hilbert space spanned by the first three Hermite-Gaussian temporal modes. The process matrix of the Raman memory interaction has 162 free parameters in three dimensions and we require at least as many measurements to fully characterise the matrix. We measure the storage efficiency for a large set of signal and control pulses constructed from random linear combinations of the first three Hermite-Gaussians (HG). We generate $k = 256$ such pulses of the form:

$$E_k(t) = \exp\left(\mathrm{i}\phi_0\right) \sum_{n=0}^{2} c_{k,n} f_n(t), \tag{6.6}$$

where $f_n(t)$ is the nth order Hermite-Gaussian polynomial, and $c_{k,n}$ are a set of random, complex coefficients generated via:

$$\begin{aligned} c_{k,0} &= \cos(\theta_1)\cos(\theta_2), \\ c_{k,1} &= \cos(\theta_1)\sin(\theta_2)\exp(\mathrm{i}\phi_1), \\ c_{k,2} &= \sin(\theta_1)\exp(\mathrm{i}\phi_2), \end{aligned} \tag{6.7}$$

and θ_i, $\phi_i \in [0, 2\pi)$.
For each pulse $E_k(t) \equiv A_k(t)\exp(\mathrm{i}\varphi_k(t))$, we generate voltage signals of the form

$$\begin{aligned} V_{1,k}(t)/V_\pi &= \arcsin\left[A_k(t)\right] + \varphi_k(t), \\ V_{2,k}(t)/V_\pi &= \arcsin\left[A_k(t)\right] - \varphi_k(t), \end{aligned}$$

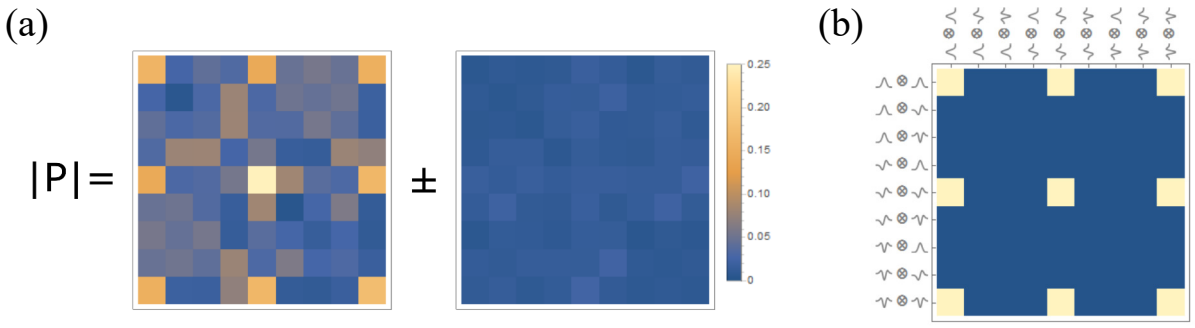


Figure 6.7: (a) Reconstructed process matrix of the Raman memory compared to (b) the ideal process matrix for a true single mode memory.

which are amplified by the RF amplifiers to drive the FMZI, as described in Section 6.2.1.

To consistently probe the same process matrix with every measurement, we need the energy of every control pulse to be the same. The coefficients $c_{k,n}$ are normalised such that the integrated intensity of all the pulses should be equal. However, slight discrepancies between the two RF amplifiers lead to a significant variation in the optical intensity of the generated pulses. We therefore measure the control pulse energy for each generated pulse and adjust it to be 150 pJ for every measurement of the storage efficiency. We also do some post processing by summing the total counts measured by the detectors in the integrated time window of the control pulses, and discarding any measurements where this varies significantly from the mean.

To reconstruct the process matrix via tomography it is crucial that we know the optical states that we use to probe the interaction. We measure the amplified voltage signals that are applied to the FMZI to generate each of the pulses using a fast oscilloscope, and using Equation 6.3 we can then infer what the optical signals are. The errors on these optical states are estimated via Monte-Carlo over small variations in the amplitudes of the two amplifiers and the DC bias voltage V_π.

The reconstructed process matrix is shown in Figure 6.7(a), and the equivalent matrix for an ideal single-mode memory is shown in Figure 6.7(b). For a given control pulse, we can extract a metric of *single-modeness* of the memory by projecting $\mathbf{P}$ onto this control mode:

$$\boldsymbol{P}_{\boldsymbol{\Omega}} = (\mathbb{1} \otimes \Omega^*)^\dagger \boldsymbol{P} (\mathbb{1} \otimes \Omega^*) \tag{6.8}$$

and finding the eigenvalues σ of the resultant matrix. We define the *single-modeness*, κ, as the largest eigenvalue of this process matrix, normalised by the sum of all the eigenvalues:

$$\kappa(\boldsymbol{\Omega}) = \frac{\max\limits_{p \in \sigma[\boldsymbol{P}_{\boldsymbol{\Omega}}]} (p)}{\sum\limits_{p \in \sigma[\boldsymbol{P}_{\boldsymbol{\Omega}}]} p}. \tag{6.9}$$

and for an ideal single-mode process, $\kappa = 1$. By projecting $\boldsymbol{P}$ onto a large set of 500,000 random control pulses from within the three-dimensional HG basis, we can extract the average, best-case and worst-case single-modeness of the reconstructed process matrix. The results are

SM	$\kappa_{\mathtt{max}}$	$\kappa_{\mathtt{avg}}$	$\kappa_{\mathtt{min}}$
κ	0.98	0.91	0.69

Table 6.1: Extracted mean and limiting cases of the single-modeness for a projection of $\boldsymbol{P}$ onto controls in a 3-dimensional subspace.

shown in Table 6.1, and we see that the memory is highly mode selective, with an average single-modeness of more that 90%. These results are very similar to those obtained from the measurement tomography of the QPG, with $\kappa = 0.920 \pm 0.024$ [227].

There are some limitations to this process tomography technique, which arise from our lack of perfect knowledge of the optical probe states, and the variation in control pulse energy between measurements. This leads to some non-physicality of the reconstructed matrix, and large errors in κ. We discuss methods to minimise these problems and improve the tomography reconstruction in Section 7.2.4.

For this reconstruction, the control pulse energy was kept to a moderate value of 150 pJ to give a storage efficiency of $\approx$ 20% (when the signal and control pulses are in the same temporal mode), and ensure that we were in the low-coupling regime with $C_s \approx 0.5$. As the coupling strength increases the memory will become more multimode, and the eigenmodes will become distorted and no longer have significant overlap with the control mode that drives the interaction. This can be thought of as an artefact of time-ordering effects [228–230] since the mapping between the fields S and B depends on the control field $\Omega(t)$ and is time dependent, and therefore the matrix describing the evolution of these coupled fields does not commute at different times. When the coupling is low, the converted field is small, and the storage interaction depends only on the overlap between the signal and control modes which propagate through the ensemble together. The spin-wave mode is uniform across the atomic ensemble and the interaction is single mode. As the coupling strength increases, the storage efficiency is higher and the magnitude of the spin-wave mode is much larger and therefore there is back-action between the stored mode and the signal, and the memory interaction depends on a three-mode overlap. This causes the interaction to become multimode, and also distorts the eigenmodes of the system. The coupling between the signal and control is large at the beginning of the atomic ensemble so the signal is predominantly read-in towards the start of the medium, and therefore the spin-wave mode is no longer spatially uniform. We explore this in more detail in Section 6.4.1.

As the coupling strength of the Raman memory increases, the process becomes less single mode, and additional eigenmodes will have a non negligible interaction with the system. This is a general property of the "broadband beam splitter equations of motion", which is also seen in the quantum pulse gate [231], and there is an inherent compromise between the single-modeness, κ, and the efficiency, η, of the device. However, it has been shown that it is possible to use cascaded QPGs, where each one individually has a low efficiency and high selectivity, to achieve a high overall efficiency and maintain good mode selectivity [225].

In an analogous way, the cavity-enhanced Raman memory will offer this capability. We discussed in Section 4.1 that the cavity-enhanced Raman memory can be seen as a perfectly mode selective device, even for high coupling strengths. We can consider this by "unfolding" the

cavity such that it is a series of free space memory interactions for each round trip of the cavity. Every individual memory has a low coupling strength and low efficiency and therefore remains highly single mode, and the spin-wave is spatially uniform throughout the medium. However, since there are multiple interactions, the interaction coherently builds up and the total efficiency can be high. In a similar way to the cascaded QPGs, the individual round trips of the cavity have low efficiency and high mode selectivity, but the process as a whole can achieve high κ and high η [205].

We have demonstrated that in the low-coupling regime the Raman memory is highly mode-selective and will map a single temporal mode into a stored spin-wave. The temporal mode of the control determines the principal eigenmode of the memory interaction, and therefore by tailoring the control pulse we can choose this eigenmode to coincide with the input signal we wish to store. All modes orthogonal to the principal eigenmode will be virtually unaffected, and transmit straight through the atomic ensemble unchanged.

6.4 Mode Manipulation

In Section 6.3 we showed that the Raman memory is highly single-mode, and therefore can be used to separate overlapping TMs – a key requirement for a TM selective device. We next investigate how the Raman memory can be used to convert between TMs.

We demonstrated above that the Raman memory stores a single temporal mode, and that this operational mode is determined by the read-in control pulse. Similarly, the optical mode that is retrieved from the memory is fully determined by the read-out control pulse and hence we can use this system to re-shape temporal modes. By selecting different modes for the read-in and read-out control pulses, we can store an optical signal in one TM and retrieve it in a different TM at a later time. As described in Section 6.1, this ability to convert temporal modes is a key element to enable qudit rotations on TMs that form a high-dimensional basis for quantum information processing, and for optimally interfacing different systems at distinct nodes of a quantum network. In this section we demonstrate conversion between Hermite-Gaussian temporal modes, and change the bandwidth of Gaussian pulses of light.

6.4.1 High Dimensional Temporal Modes

We discussed in Section 6.1.1 that temporal modes are a promising platform for encoding photonic information, as they span an infinite dimensional Hilbert space and are compatible with single mode fibres. As an example we investigate using Hermite-Gaussian (HG) polynomials as a qudit basis for photons, although any orthonormal basis could be used. Here we demonstrate converting between basis states in this high dimensional basis, by generating and storing an optical signal in an nth order mode, HG_n, and retrieving it in an mth order mode, HG_m. Two examples of HG mode conversion are shown in Figure 6.8(a), where we demonstrate re-shaping of HG_0 to HG_2 and re-shaping of HG_1 to HG_3, respectively.

The efficiency of the Raman memory interaction is dependent on the integrated control Rabi frequency, and not on the temporal mode shape or bandwidth of the control pulse [192].

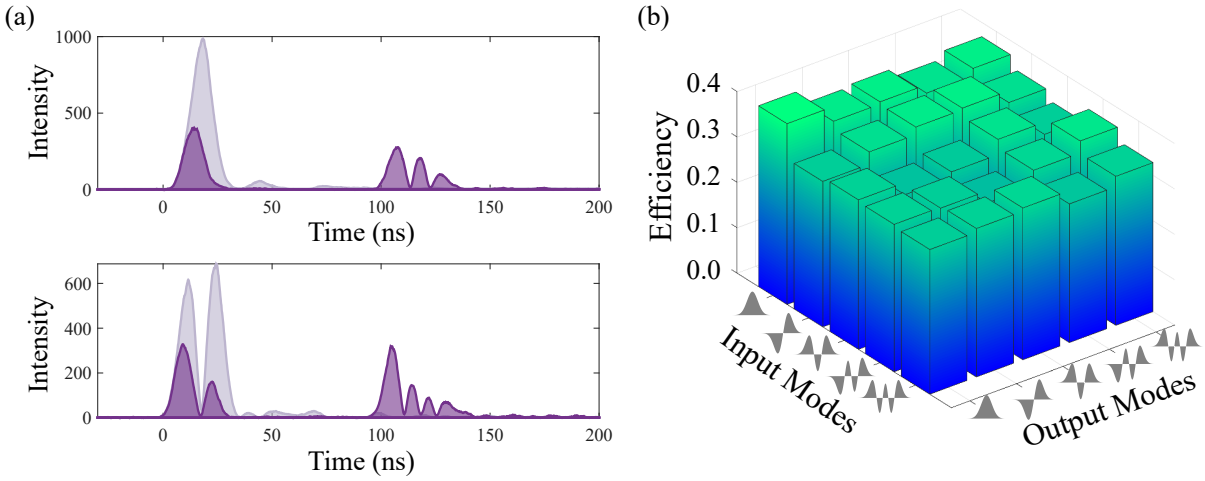


Figure 6.8: (a) Temporal mode conversion between Hermite-Gaussian modes using the Raman memory. Upper: conversion from HG_0 to HG_2. Lower: conversion from HG_1 to HG_3. (b) The efficiency of conversion between the first five Hermite-Gaussian temporal modes.

Therefore, if the control pulse energy is fixed, the efficiency of the Raman memory interaction will be the same, regardless of the mode of the storage and retrieval pulses. Figure 6.8(b) shows the efficiency of mode conversion between the first five HG basis states where the control pulse energy is kept constant, and we see that the efficiency of HG mode conversion is approximately constant at $\eta_{\text{mem}} \approx 35\%$ for all basis state transformations.

We note in Figure 6.8(a) that the output pulse shape is distorted and has a low overlap with the ideal HG basis states. We investigate this as a function of the control pulse energy, and find that the pulse becomes more and more distorted as the Raman coupling strength increases, as shown in Figure 6.9. This is a consequence of the broadband beam splitter nature of the equations of motion that govern the Raman memory that was discussed in Section 6.3. As the coupling strength increases, the eigenmodes of the Raman memory change and no longer have a large overlap with the temporal mode of the control pulse that drives the memory. This is well captured by our numerical simulations of the system, as shown in the lower panel of Figure 6.9. As we discuss below in Section 6.5, we can compensate for this by adapting the temporal mode of the retrieval control pulse such that the principal eigenmode of the output interaction corresponds to the desired temporal mode in the HG basis.

We also notice in Figure 6.9 that the pulse that is not stored in the memory but transmitted straight through appears temporally offset from the input pulse, which is not present in our simulations. This is because there is a slight timing mismatch between the signal and control pulses for these data. The control pulse arrives at the memory slightly later than the signal and therefore the later part of the signal pulse is read-in to the memory more efficiently than the leading part. This could be compensated for by adjusting the relative time delay between the signal and control pulses, and would also increase the storage efficiency.

6.4.2 Bandwidth Conversion

Bandwidth conversion has been previously demonstrated in a Raman memory in diamond [232], and photons with a 4.1 nm bandwidth were converted by a factor of two in both directions. More

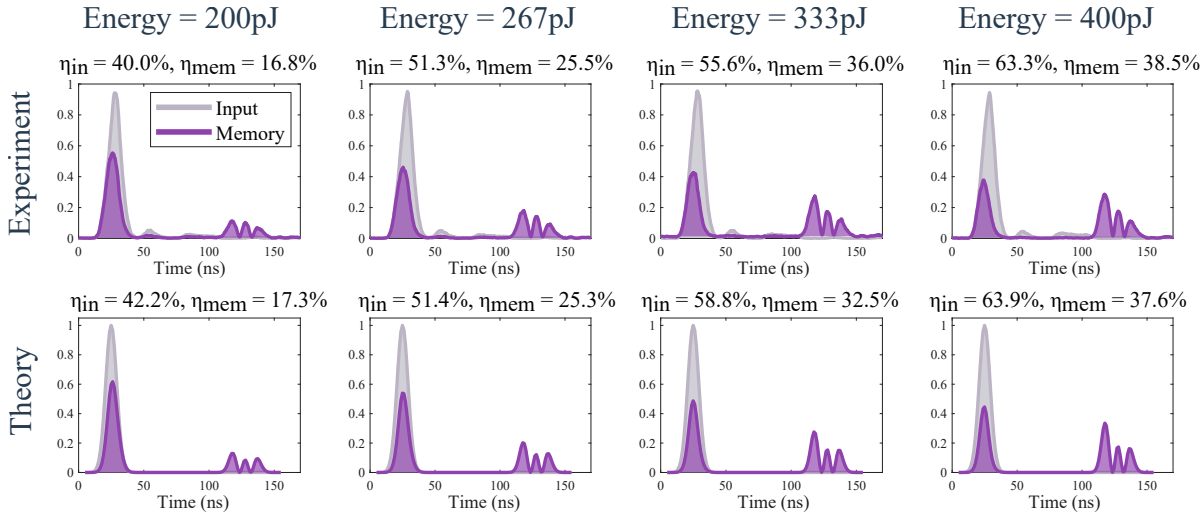


Figure 6.9: Mode conversion from a HG_0 to HG_2 as a function of control pulse energy. The upper panel shows the experimental data, and we see significant pulse distortion from the ideal Hermite-Gaussian basis as the coupling strength increases. The lower panel shows the numerical simulation of the system, and we see very strong agreement with the experiment. The model captures the output shape of the memory very well, and also agrees strongly with the numerical values of the read in efficiency (η_{in}), and total efficiency (η_{mem}).

recently, bandwidth conversion was performed in an ATS memory [174], and narrowband pulses were compressed and stretched by a factor of two. Here we investigate bandwidth conversion in the MHz–GHz regime which can be used to interface narrowband atomic systems with fast GHz-bandwidth communication networks by converting the bandwidth of pulses between these two regimes, and we push the bandwidth conversion beyond a factor of two.

We demonstrate bandwidth conversion of ns-duration Gaussian pulses using the Raman memory by generating Gaussian read-in and read-out control pulses of different durations. We choose the read-in control pulse to have the same pulse duration as the input signal to maximise the storage efficiency. The data in Figure 6.10 shows we can both increase and decrease the bandwidth of the pulses by at least a factor of 10.

We adjust the amplitudes of the pulses such that the pulse energy is the same, independent of pulse duration. We normalise the amplitude with respect to the shortest duration pulse, and hence the pulse energy is equal but low for all pulses. Nevertheless, we see that the memory efficiency for bandwidth conversion by a factor of 4 in both directions is approximately constant if the overall control pulse energy is the same, as shown in Figure 6.11(a).

We also generate pulses for bandwidth conversion that have the same peak amplitude and therefore the longer duration pulses have higher pulse energy. The efficiency of the conversion to longer pulses is much more efficient since the retrieval interaction strength is higher, as evident in Figure 6.11(b). We compare the efficiency of this process to that of using a passive bandwidth filter by calculating the transmission of a Gaussian pulse through a top-hat frequency filter. We note that a passive filter can only decrease the bandwidth, unlike the Raman memory which can convert in both directions. The efficiency of filtering using a top-hat filter decreases quickly with the compression factor, as the majority of the pulse energy is lost. At large bandwidth compression factors the memory can outperform a passive filter.

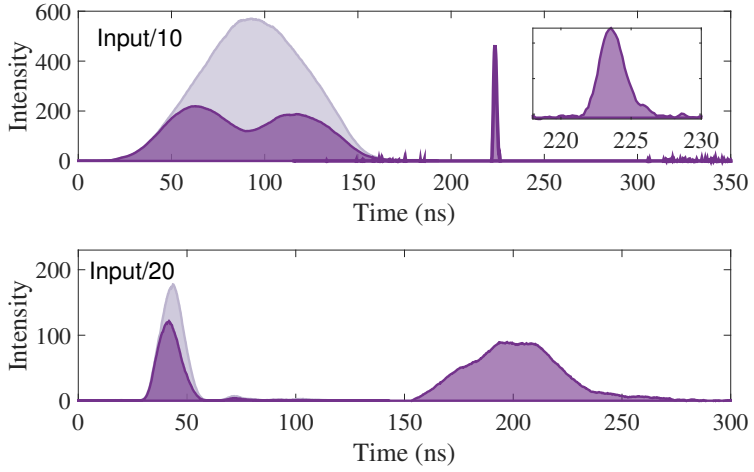


Figure 6.10: Bandwidth conversion using a Raman memory to both decrease and increase the pulse duration. Upper: the input pulse has a full-width half-maximum of 50 ns, and is converted to an output pulse with 1 ns duration. The inset shows a zoom in of the output pulse. Lower: an input pulse with duration 10 ns is converted to an output pulse with duration 100 ns.

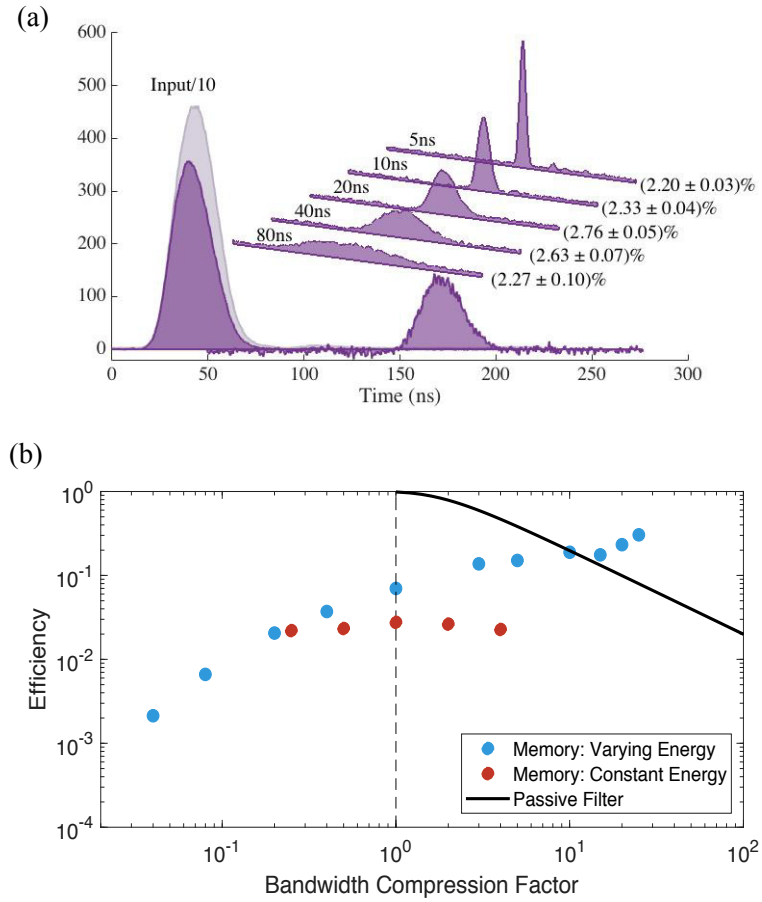


Figure 6.11: (a) Bandwidth conversion for a constant control pulse energy. The input pulse has duration 20 ns, and the inset shows the retrieval time bin when the output pulse ranges from 5 ns to 80 ns. (b) The efficiency of bandwidth conversion using the Raman memory as a function of the bandwidth compression. Higher compression factors indicate conversion to longer pulse durations. The solid line shows the equivalent efficiency for a passive top-hat bandwidth filter.

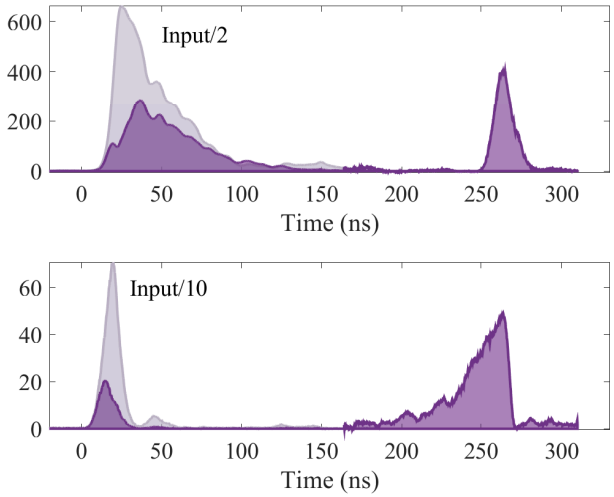


Figure 6.12: Upper: Mode conversion from a falling exponential to a Gaussian temporal mode. The read-in efficiency is $\eta_{\text{in}} \approx 48\%$, and the total conversion efficiency $\eta_{\text{mem}} \approx 12\%$. Lower: Mode conversion from a Gaussian to a rising exponential temporal mode, with $\eta_{\text{in}} \approx 69\%$ and $\eta_{\text{mem}} \approx 13\%$. The oscillations in the exponential waveforms are due to the finite sampling of the AWGs which leads to ringing of the RF amplifiers due to the abrupt changes in amplitude.

This proof of principal experiment demonstrates that we can perform significant bandwidth conversion by up to a factor of 25 in both directions, and that if we keep the pulse energy constant the efficiency is independent of the bandwidth conversion factor.

The ability to tailor the temporal mode of the control pulse can enable us to optimise light-matter interactions. For example, to optimally couple different devices and components of a quantum network we require the capability to engineer and manipulate the temporal mode of photons, tailoring them for maximum coupling to distinct technologies at different nodes. In Figure 6.12 we demonstrate that we can convert between a falling exponential, which is the temporal mode emitted by a two-level atomic system, and HG modes which closely resemble the eigenmodes for photon pair sources [233]. We also demonstrate conversion in the opposite direction: from a Gaussian to a rising exponential which is optimally absorbed by a two-level system. Whilst the efficiency of conversion is low, this demonstrates how the Raman memory could improve coupling of light-matter systems.

6.5 Optimising the Raman Light-Matter Interaction

We have demonstrated that the Raman memory is a mode selective device that acts on a single, user-defined TM, and can be used to convert between different TMs. This establishes the Raman memory as a powerful new tool for TM operation. As a final application, we consider how tailoring the temporal mode of light can directly enhance the Raman light-matter interaction.

To see this we consider the mapping of the input signal mode to the spin-wave which is described by a storage Green's function K_1, which can be known either from a numerical

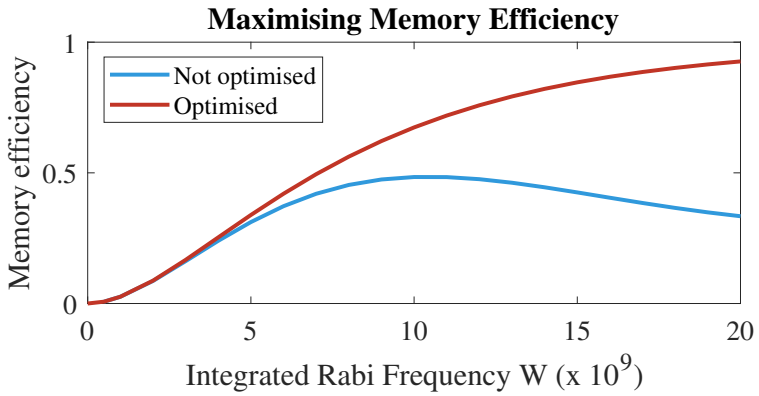


Figure 6.13: Blue: the memory efficiency for a Gaussian signal read-in by a Gaussian control pulse as a function of the integrated Rabi frequency of the control field. Red: the memory efficiency of a Gaussian signal mode using the optimised control pulse. Much higher efficiencies can be reached by compensating for the AC Stark shift, and matching the optimal mode of the system.

simulation, an analytic approximation, or a tomographical reconstruction as demonstrated in Section 6.3. From this Green's function we can find the principal input eigenmode, that is the signal mode that will be optimally stored in the memory for any control temporal mode. The principal eigenmode of the memory interaction depends on the integrated Rabi frequency of the control field and its temporal profile, and as we change $\Omega_{\text{in}}(t)$ it is straightforward to calculate the new optimal mode by a simple linear map. We now wish to consider the reverse problem: we want to find the control temporal profile $\Omega_{\text{in}}(t)$ such that the optimal mode matches our pre-determined input signal. Similarly, for the retrieval process, we want to find the optimal read-out control pulse $\Omega_{\text{out}}(t)$ such that the output mode of the memory matches a chosen temporal mode.

This inversion problem can be solved by minimising the difference between the optimal eigenmode of the system and our desired temporal mode, for example using a 4th order Runge-Kutta procedure, and this numerical method has been developed by J.H.D.M. and J.K. Figure 6.13 shows the predicted storage efficiency of a given input signal mode as a function of the control Rabi frequency using, first, the same control TM as the signal (non-optimised), and secondly the optimised control pulse such that the input signal matches the eigenmode of the system. We can see that at high control pulse energies the storage efficiency for the optimised control mode is significantly higher than for the non-optimised case and there is no longer a turn over in the storage efficiency from the AC Stark shift. This is because the optimised mode compensates for the AC Stark shift by adding a time varying phase to the control pulse such that the signal and control fields remain in two-photon resonance throughout the interaction. We can therefore combat the AC Stark shift by shaping the temporal mode of either the signal or control pulse, and reach much higher memory efficiency.

As a second example, we consider the mode conversion between Hermite-Gaussians presented in Figure 6.9, where the output mode becomes distorted at high coupling strengths. By adjusting the retrieval control pulse we can see, numerically, that we can alter the output mode such that it has much greater overlap with the ideal HG_2 polynomial, as shown in Figure 6.14. This optimisation procedure has not yet been investigated experimentally, but highlights the exciting

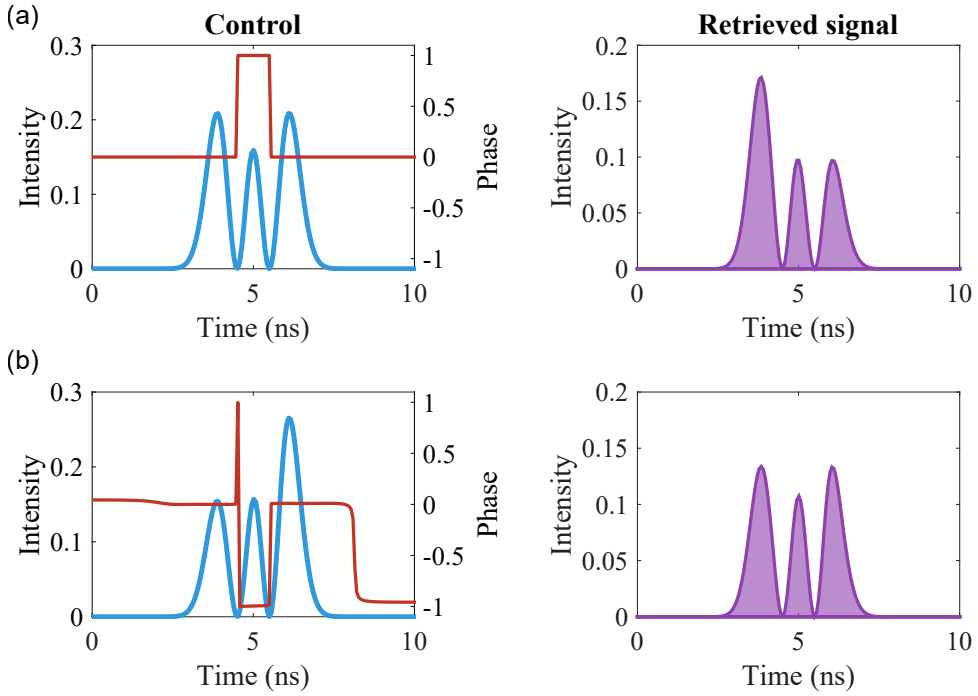


Figure 6.14: Simulation of optimising the output temporal mode of the memory by tailoring the amplitude (blue) and phase (red, in units of π) of the control mode. Upper: Simulations of the output mode (purple) when the retrieval control pulse is a second-order Hermite-Gaussian with an energy of 350 pJ. Lower: by tailoring the control pulse temporal mode but keeping the energy the same, we can retrieve a signal that has high overlap with the ideal Hermite-Gaussian basis state.

future possibilities of the Raman memory to enhance and optimise light-matter interactions. We are hopeful that by tailoring the control pulses that drive the Raman memory we will soon be able to reach significantly higher memory efficiencies, and convert between user-defined temporal modes with high fidelity.

6.6 Summary

In this chapter I have demonstrated that the Raman memory is a powerful new platform for operating on temporal modes of light, and the key results are outlined below.

- In Section 6.1 I introduced temporal modes as an appealing basis for photonic quantum information processing, and highlighted the application of tailoring the temporal mode of light to enhance light-matter interactions. I introduced the requirements for such a device to operate on temporal modes.
- In Section 6.2 I described the experimental set-up that we built to investigate the modal properties of the Raman memory.
- In Section 6.3 I presented a process tomography of the storage interaction of the Raman memory, and demonstrated that, in the low-coupling regime, the Raman memory is effectively single mode and can be used to address an individual temporal mode of light. By adjusting the temporal mode of the control pulse that drives the memory we can determine the principal eigenmode of the memory that will be optimally stored.

- In Section 6.4 I demonstrated mode manipulation between temporal modes of light. By selecting different modes for the read-in and read-out control pulses we can store an input signal in one temporal mode and retrieve it in another. I used this system to demonstrate conversion between Hermite-Gaussian polynomials and significantly alter the bandwidth of light pulses as well as conversion to and from an exponential waveform that would optimally interface with a two-level system.
- In Section 6.5 I investigated applications of the Raman memory to enhance light-matter interactions. I discussed how we can tailor the control pulses that drive the Raman memory to maximise the memory efficiency, and match the principal eigenmode of the system to one of our choice.

This highlights the capability of the Raman memory to separate and manipulate temporal modes of light for a significant range of applications. We have developed a method to reconstruct the underlying process matrix of an atomic memory and benchmark the mode selectivity. We have also built infrastructure to prepare, separate, and manipulate temporal modes of light which will enable many exciting future experiments outlined in Section 7.2.

Chapter 7:

Conclusions and Outlook

Quantum memories have been highlighted as a vital component of photonic networks, and can be used to synchronise probabilistic processes in a temporal multiplexing scheme. In this thesis we have focused on the Raman quantum memory which can store broadband signals and has a large synchronising capacity, and is therefore ideally suited for local synchronisation applications. I have investigated methods to increase the signal to noise ratio of the Raman memory, and also explored the temporal mode properties of the memory. The key findings are summarised below.

7.1 Summary

In the first half of this thesis I investigated methods to increase the signal to noise ratio (SNR) of the Raman memory, by maximising the memory efficiency and by suppressing the unwanted noise processes.

In Chapter 3 I investigated methods to increase the Raman coupling constant by increasing the operational optical depth of the memory. By introducing a molecular buffer gas to the warm caesium vapour I was able to improve the spin polarisation of the ensemble at high temperatures and increase the optical depth by a factor of four. I demonstrated dramatically higher Raman coupling in the system at high temperatures and, even though the system supports four-wave mixing gain, we can predict that the true memory efficiency is significantly higher.

In Chapter 4 I investigated a method to suppress four-wave mixing noise by operating the memory inside a low finesse optical cavity to suppress anti-Stokes scattering. The cavity successfully reduced the relative size of the four-wave mixing noise process with respect to the desired memory interaction, and enhanced the internal signal to noise ratio by an order of magnitude. However, intra-cavity loss significantly reduces the memory efficiency, and it is experimentally challenging to reduce this to the level where we would be able to observe quantum storage.

In Chapter 5 I investigated an alternative method of suppressing four-wave mixing noise by introducing an absorption feature at the frequency of the unwanted anti-Stokes field. By operating the memory at a specific detuning from resonance we can ensure that the anti-Stokes field is on resonance with the atomic transition, and dramatically inhibit anti-Stokes scattering. I observed an increase in the signal to noise ratio by an order of magnitude using this technically simple noise suppression scheme. We predict that the residual noise is a combination of spontaneous scattering due to imperfect optical pumping fluorescence noise, and discussed methods to decrease this. Nevertheless, this built-in noise suppression scheme has significantly reduced the amount of noise in the Raman memory protocol, to the level where we predict non-classical recall with a single-photon source with a heralding efficiency greater than 10.5%. This scheme is applicable to any off-resonant memory protocol, and will enable noise suppression across a wide variety of memory protocols and platforms.

In Chapter 6 I investigated the modal properties of the Raman memory. By performing a full process tomography on the memory interaction we quantified the degree to which the memory is single mode, and found an average single-modeness of over 90% in the Hilbert space spanned by the first three Hermite-Gaussian modes. The memory is hence effectively single-mode in this low-coupling regime ($C_s \lesssim 1$) and can be used as a device to separate and manipulate temporal modes of light. I demonstrated mode conversion between different Hermite-Gaussian modes, and also bandwidth conversion of Gaussian pulses which paves the way towards using this light-matter interface as a device for temporal mode manipulation in photonic quantum networks.

7.2 Future Work

There are many directions in which we can build on the results presented in this thesis, and I outline a few of these avenues below.

7.2.1 Further Noise Suppression

The results presented in Chapter 5 demonstrate the power of the built-in noise suppression scheme to reduce four-wave mixing noise. We believe that the remaining noise is a combination of fluorescence noise due to linear absorption of the control field, and spontaneous Raman scattering due to imperfect optical pumping. We discussed in Section 5.4 that the frequency filtering after the memory (described in detail in Section 3.1.3) is not optimal for this operational detuning. Four of the etalons have a free-spectral range of 18.4 GHz and therefore will transmit the majority of the fluorescence noise. We intend to change the filtering set-up to use etalons with a different free-spectral range to enable larger suppression of the fluorescence noise. We also discussed in Section 5.4 that the extinction of the EOM that switches off the optical pumping during storage has significantly degraded, and this decreases the memory efficiency and increases the noise due to spontaneous Raman scattering. We will replace this with a high extinction Pockels cell to achieve high optical pumping efficiency and low spin-wave depletion during storage. We hope that the combination of these two experimental changes will significantly reduce the noise level and enable storage and retrieval of single photons in the near future.

7.2.2 Temporal Multiplexing

If the noise level can be reduced further then we will be able to demonstrate temporal multiplexing of probabilistic single photon sources. As a first demonstration, we will store a heralded single photon in the memory using the source outlined in Section 4.4. Once we detect a subsequent herald event from the source we will retrieve the stored photon from the memory and temporally synchronise it with the second photon from the source in a two-photon state. This will decrease the waiting time for a two-photon state, compared to waiting for two separate sources to generate a heralded single photon at the same time. Due to the large synchronising capacity of the Raman memory, we hope to be able to demonstrate synchronisation of multi-photon states using either multiple memories and a single source, or multiple memories interfaced with multiple sources.

7.2.3 Light-Matter HOM

We described in Section 2.2.3 that the Raman memory acts as a light-matter beam splitter between the signal mode and the stored spin-wave. We want to demonstrate this by performing Hong-Ou-Mandel type interference between these two modes, in a similar way to the experiment performed in [234]. To do this we would need to tune the read-in and read-out efficiency of the Raman memory to 50% and then store a single photon in the memory. If a second photon is sent to the memory then Hong-Ou-Mandel interference dictates that there will be bunching between the bosonic photon and spin-wave modes. Either the first photon will be retrieved from the memory and the second photon will be transmitted through the memory, or the first photon will not be retrieved and the second photon will be stored, i.e. we will generate a superposition of two photons or two spin-waves in the state $(|0_S 2_B\rangle + |2_S 0_B\rangle)/\sqrt{2}$. The $|1_S 1_B\rangle$ terms cancel out and we should not have any incidences of a photon being detected in both the first and second time bins, analogous to Hong-Ou-Mandel interference between photons at a beam splitter.

7.2.4 Temporal Mode Verification

The results on the temporal mode properties of the Raman memory presented in Chapter 6 are in their infancy and there are many ways we can develop this line of research. Firstly, we want to improve the reliability of the process tomography method by having better knowledge of the optical probe states used to perform the reconstruction. A technique has been developed by J.H.D.M. to characterise pulses of a few nanosecond duration, and is analogous to SPIDER which characterises ultrafast pulses [235]. This technique, known as Temporal Electric-field Amplitude and Phase: Algorithmic Reconstruction via Time-domain Interferometry (TEAPARTI), allows single-shot characterisation of the amplitude and phase of an electric field, and more details are given in [190]. We are in the process of benchmarking this technique, and thus far have seen that it gives a much more accurate prediction of an optical state than we get from measuring the electrode signals. We therefore intend to redo the process tomography, and measure the optical signal and control pulses via TEAPARTI in parallel to measuring the storage efficiency in the memory. This should enable a robust reconstruction of the memory process matrix, that is insensitive to experimental drifts and imperfections.

We will also be able to use TEAPARTI to verify temporal mode manipulation. The results presented in Section 6.4 on conversion between different temporal modes are thus far qualitative only, as we only had access to the optical intensity measured by an APD and recorded by a timetagger. By using TEAPARTI we will be able to reconstruct the amplitude and phase of the retrieved pulse, and will be able to quantify the overlap between this and the desired output mode.

7.2.5 High-Efficiency Single-Mode Cavity Memory

We discussed in Section 6.5 that we can tailor the temporal mode of the control pulse to increase the memory efficiency. We can choose the mode of the read-in pulse such that the principal eigenmode of the storage interaction matches the input signal mode and compensates for the AC Stark shift to achieve high efficiency coupling. We can choose the read-out control pulse to retrieve the spin-wave efficiently into a different temporal mode of our choice. We hope to be able to demonstrate efficiencies approaching unity in the Raman memory by appropriate shaping of the control pulses.

As the efficiency or coupling strength of the Raman memory increases the memory will become more multimode, as described in Section 6.3. However, the cavity memory remains single mode even in the high coupling regime. By operating the memory inside the cavity we will be able to reach high single-modeness and high efficiency simultaneously. This demonstrates a powerful advantage of the cavity-enhanced Raman memory. In addition to suppressing four-wave mixing noise, it allows for a high efficiency, highly single-mode interaction. We therefore hope to explore the temporal mode properties of the cavity-enhanced Raman memory. We can perform a process tomography on the cavity-enhanced Raman memory and quantify how single mode the system is, and we then will be able to demonstrate high-efficiency temporal mode selectivity and manipulation.

Furthermore, we can combine the BNS scheme for suppressing four-wave mixing and the cavity to achieve a memory with high mode selectivity and excellent noise suppression. Since the cavity enhances the memory interaction it will be possible to operate the memory with a lower control pulse energy which will also reduce the amount of fluorescence noise. By combining the different techniques presented in this thesis it will be possible to build a memory with high efficiency, high mode-selectivity and low noise.

7.3 Conclusions

In this thesis I have increased the signal-to-noise ratio in the Raman memory by an order of magnitude, and demonstrated that the Raman memory is a useful device for temporal mode manipulation. I have highlighted routes to improve the signal-to-noise ratio even further, and discussed how the cavity-enhanced Raman memory will enable high-efficiency, single-mode operation. The results in this thesis and the subsequent discussions represent key stepping stones towards an efficient, low-noise, mode-selective quantum memory.

Bibliography

[1] V. Giovannetti, S. Lloyd, and L. Maccone, *Quantum-Enhanced Measurements: Beating the Standard Quantum Limit*, Science **306**, 1330 (2004).

[2] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing, Bangalore, India , 175 (1984).

[3] R. P. Feynman, *Simulating physics with computers*, International Journal of Theoretical Physics **21**, 467 (1982).

[4] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, *Quantum computers*, Nature **464**, 45 (2010), Review Article.

[5] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, *Elementary gates for quantum computation*, Physical Review A **52**, 3457 (1995).

[6] D. P. DiVincenzo, *Two-bit gates are universal for quantum computation*, Physical Review A **51**, 1015 (1995).

[7] P. W. Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, SIAM Journal on Computing **26**, 1484 (1997).

[8] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, pp. 212–219, New York, NY, USA, 1996.

[9] D. Deutsch and R. Jozsa, *Rapid solution of problems by quantum computation*, Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **439**, 1907 (1992).

[10] R. Raussendorf and H. J. Briegel, *A One-Way Quantum Computer*, Physical Review Letters **86**, 5188 (2001).

[11] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *Quantum Computation by Adiabatic Evolution*, `arXiv` quant-ph/0001106 (2000).

[12] R. Babbush, P. J. Love, and A. Aspuru-Guzik, *Adiabatic Quantum Simulation of Quantum Chemistry*, Scientific Reports **4**, 6603 (2014).

[13] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Simulated Quantum Computation of Molecular Energies*, Science **309**, 1704 (2005).

[14] S. A. Cook, *An Overview of Computational Complexity*, Communications of the ACM **26**, 400 (1983).

[15] P. Date, R. Patton, C. Schuman, and T. Potok, *Efficiently embedding QUBO problems on adiabatic quantum computers*, Quantum Information Processing **18**, 117 (2019).

[16] D. P. Nazareth and J. D. Spaans, *First application of quantum annealing to IMRT beamlet intensity optimization*, Physics in Medicine and Biology **60**, 4137 (2015).

[17] S. Adachi and M. Henderson, *Application of Quantum Annealing to Training of Deep Neural Networks*, `arXiv` 1510.06356 [quant-ph] (2015).

[18] A. D. King *et al.*, *Observation of topological phenomena in a programmable lattice of 1,800 qubits*, Nature **560**, 456 (2018).

[19] W. K. Wootters and W. H. Zurek, *A single quantum cannot be cloned*, Nature **299**, 802 (1982).

[20] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels*, Physical Review Letters **70**, 1895 (1993).

[21] A. K. Ekert, *Quantum cryptography based on Bell's theorem*, Physical Review Letters **67**, 661 (1991).

[22] V. Giovannetti, S. Lloyd, and L. Maccone, *Advances in quantum metrology*, Nature Photonics **5**, 222 (2011), Review Article.

[23] G. Tóth and I. Apellaniz, *Quantum metrology from a quantum information science perspective*, Journal of Physics A: Mathematical and Theoretical **47**, 424006 (2014).

[24] M. Napolitano, M. Koschorreck, B. Dubost, N. Behbood, R. J. Sewell, and M. W. Mitchell, *Interaction-based quantum metrology showing scaling beyond the Heisenberg limit*, Nature **471**, 486 (2011).

[25] S. Boixo, S. T. Flammia, C. M. Caves, and J. Geremia, *Generalized Limits for Single-Parameter Quantum Estimation*, Physical Review Letters **98**, 090401 (2007).

[26] D. P. DiVincenzo, *The Physical Implementation of Quantum Computation*, Fortschritte der Physik **48**, 771 (2000).

[27] H. Weimer, M. Müller, I. Lesanovsky, P. Zoller, and H. P. Büchler, *A Rydberg quantum simulator*, Nature Physics **6**, 382 (2010).

[28] D. Kielpinkski, C. Monroe, and D. J. Wineland, *Architecture for a large-scale ion-trap quantum computer*, Nature **417**, 709 (2002).

[29] J. Clarke and F. K. Wilhelm, *Superconducting quantum bits*, Nature **453**, 1031 (2008).

[30] B. E. Kane, *A silicon-based nuclear spin quantum computer*, Nature **393**, 133 (1998).

[31] D. D. Awschalom, R. Hanson, J. Wrachtrup, and B. B. Zhou, *Quantum technologies with optically interfaced solid-state spins*, Nature Photonics **12**, 516 (2018).

[32] J. L. O'Brien, A. Furusawa, and J. Vuckovic, *Photonic quantum technologies*, Nature Photonics **3**, 687 (2009).

[33] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, *High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits*, Physical Review Letters **117**, 060504 (2016).

[34] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, *Experimental comparison of two quantum computing architectures*, Proceedings of the National Academy of Sciences **114**, 3305 (2017).

[35] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, *Fast quantum logic gates with trapped-ion qubits*, Nature **555**, 75 (2018).

[36] N. Friis *et al.*, *Observation of Entangled States of a Fully Controlled 20-Qubit System*, Physical Review X **8**, 021012 (2018).

[37] C. Hempel *et al.*, *Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator*, Physical Review X **8**, 031022 (2018).

[38] I. Bloch, *Ultracold quantum gases in optical lattices*, Nature Physics **1**, 23 (2005).

[39] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, *Fast Quantum Gates for Neutral Atoms*, Physical Review Letters **85**, 2208 (2000).

[40] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman, *Demonstration of a Neutral Atom Controlled-NOT Quantum Gate*, Physical Review Letters **104**, 010503 (2010).

[41] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys, *Entanglement of Two Individual Neutral Atoms Using Rydberg Blockade*, Physical Review Letters **104**, 010502 (2010).

[42] Y. Wang, A. Kumar, T.-Y. Wu, and D. S. Weiss, *Single-qubit gates based on targeted phase shifts in a 3D neutral atom array*, Science **352**, 1562 (2016).

[43] Y.-Y. Jau, A. M. Hankin, T. Keating, I. H. Deutsch, and G. W. Biedermann, *Entangling atomic spins with a Rydberg-dressed spin-flip blockade*, Nature Physics **12**, 71 (2015).

[44] Y. O. Dudin, L. Li, and A. Kuzmich, *Light storage on the time scale of a minute*, Physical Review A **87**, 031801 (2013).

[45] H. Bernien *et al.*, *Probing many-body dynamics on a 51-atom quantum simulator*, Nature **551**, 579 (2017).

[46] G. Wendin, *Quantum information processing with superconducting circuits: a review*, Reports on Progress in Physics **80**, 106001 (2017).

[47] R. Barends *et al.*, *Superconducting quantum circuits at the surface code threshold for fault tolerance*, Nature **508**, 500 (2014).

[48] *IBM Q - Quantum Computing*, 2018, https://www.research.ibm.com/ibm-q/.

[49] W. Knight, *IBM announces a trailblazing quantum machine*, 2017, https://www.technologyreview.com/s/609451/ibm-raises-the-bar-with-a-50-qubit-quantum-computer/.

[50] J. Hsu, *CES 2018: Intel's 49-Qubit Chip Shoots for Quantum Supremacy*, 2018, https://spectrum.ieee.org/tech-talk/computing/hardware/intels-49qubit-chip-aims-for-quantum-supremacy.

[51] *Google aims for quantum supremacy* Physics World, 2018, https://physicsworld.com/a/google-aims-for-quantum-supremacy/.

[52] *D-Wave Systems*, https://www.dwavesys.com/press-releases/d-waveannouncesd-wave-2000q-quantum-computer-and-first-system-order.

[53] J. J. L. Morton, A. M. Tyryshkin, R. M. Brown, S. Shankar, B. W. Lovett, A. Ardavan, T. Schenkel, E. E. Haller, J. W. Ager, and S. A. Lyon, *Solid-state quantum memory using the 31P nuclear spin*, Nature **455**, 1085 (2008).

[54] K. Saeedi, S. Simmons, J. Z. Salvail, P. Dluhy, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, J. J. L. Morton, and M. L. W. Thewalt, *Room-Temperature Quantum Bit Storage Exceeding 39 Minutes Using Ionized Donors in Silicon-28*, Science **342**, 830 (2013).

[55] M. Koch, J. G. Keizer, P. Pakkiam, D. Keith, M. G. House, E. Peretz, and M. Y. Simmons, *Spin read-out in atomic qubits in an all-epitaxial three-dimensional transistor*, Nature Nanotechnology **14**, 137 (2019).

[56] D. Loss and D. P. DiVincenzo, *Quantum computation with quantum dots*, Physical Review A **57**, 120 (1998).

[57] F. H. L. Koppens, C. Buizert, K. J. Tielrooij, I. T. Vink, K. C. Nowack, T. Meunier, L. P. Kouwenhoven, and L. M. K. Vandersypen, *Driven coherent oscillations of a single electron spin in a quantum dot*, Nature **442**, 766 (2006).

[58] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard, *Coherent Manipulation of Coupled Electron Spins in Semiconductor Quantum Dots*, Science **309**, 2180 (2005).

[59] P. Harvey-Collard *et al.*, *Coherent coupling between a quantum dot and a donor in silicon*, Nature Communications **8**, 1029 (2017).

[60] E. Kawakami, P. Scarlino, D. R. Ward, F. R. Braakman, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen, *Electrical control of a long-lived spin qubit in a Si/SiGe quantum dot*, Nature Nanotechnology **9**, 666 (2014).

[61] M. Veldhorst *et al.*, *A two-qubit logic gate in silicon*, Nature **526**, 410 (2015).

[62] A. P. Nizovtsev, S. Y. Kilin, F. Jelezko, T. Gaebal, I. Popa, A. Gruber, and J. Wrachtrup, *A quantum computer based on NV centers in diamond: Optically detected nutations of single electron and nuclear spins*, Optics and Spectroscopy **99**, 233 (2005).

[63] D. J. Christle, A. L. Falk, P. Andrich, P. V. Klimov, J. U. Hassan, N. Son, E. Janzén, T. Ohshima, and D. D. Awschalom, *Isolated electron spins in silicon carbide with millisecond coherence times*, Nature Materials **14**, 160 (2014).

[64] M. Zhong, M. P. Hedges, R. L. Ahlefeldt, J. G. Bartholomew, S. E. Beavan, S. M. Wittig, J. J. Longdell, and M. J. Sellars, *Optically addressable nuclear spins in a solid with a six-hour coherence time*, Nature **517**, 177 (2015).

[65] H. J. Kimble, *The quantum internet*, Nature **453**, 1023 (2008).

[66] T. Rudolph, *Why I am optimistic about the silicon-photonic route to quantum computing*, APL Photonics **2**, 030901 (2017).

[67] D. Dimitropoulos and B. Jalali, *Noise and Information Capacity in Silicon Nanophotonics*, IEEE Photonics Journal **7**, 1 (2015).

[68] A. Crespi, R. Ramponi, R. Osellame, L. Sansoni, I. Bongioanni, F. Sciarrino, G. Vallone, and P. Mataloni, *Integrated photonic quantum gates for polarization qubits*, Nature Communications **2**, 566 (2011).

[69] J. P. Lee, L. M. Wells, B. Villa, S. Kalliakos, R. M. Stevenson, D. J. P. Ellis, I. Farrer, D. A. Ritchie, A. J. Bennett, and A. J. Shields, *Controllable Photonic Time-Bin Qubits from a Quantum Dot*, Physical Review X **8**, 021078 (2018).

[70] L. Olislager, J. Cussey, A. T. Nguyen, P. Emplit, S. Massar, J.-M. Merolla, and K. P. Huy, *Frequency-bin entangled photons*, Physical Review A **82**, 013804 (2010).

[71] B. Brecht, A. Eckstein, R. Ricken, V. Quiring, H. Suche, L. Sansoni, and C. Silberhorn, *Demonstration of coherent time-frequency Schmidt mode selection using dispersion-engineered frequency conversion*, Physical Review A **90**, 030302 (2014).

[72] Z.-Y. Zhou, Y. Li, D.-S. Ding, W. Zhang, S. Shi, B.-S. Shi, and G.-C. Guo, *Orbital angular momentum photonic quantum interface*, Light: Sciences and Applications **5**, e16019 (2016).

[73] E. Knill, R. Laflamme, and G. J. Milburn, *A scheme for efficient quantum computation with linear optics*, Nature **409**, 46 (2001).

[74] T. Guerreiro, A. Martin, B. Sanguinetti, J. S. Pelc, C. Langrock, M. M. Fejer, N. Gisin, H. Zbinden, N. Sangouard, and R. T. Thew, *Nonlinear Interaction between Single Photons*, Physical Review Letters **113**, 173601 (2014).

[75] T. Peyronel, O. Firstenberg, Q.-Y. Liang, S. Hofferberth, A. V. Gorshkov, T. Pohl, M. D. Lukin, and V. Vuletic, *Quantum nonlinear optics with single photons enabled by strongly interacting atoms*, Nature **488**, 57 (2012).

[76] A. Javadi *et al.*, *Single-photon non-linear optics with a quantum dot in a waveguide*, Nature Communications **6**, 8655 (2015).

[77] R. J. Glauber, *The Quantum Theory of Optical Coherence*, Physical Review **130**, 2529 (1963).

[78] S. Tanzilli, H. De Riedmatten, H. Tittel, H. Zbinden, P. Baldi, M. De Micheli, D. B. Ostrowsky, and N. Gisin, *Highly efficient photon-pair source using periodically poled lithium niobate waveguide*, Electronics Letters **37**, 26 (2001).

[79] A. B. U'Ren, C. Silberhorn, K. Banaszek, I. A. Walmsley, R. Erdmann, W. P. Grice, and M. G. Raymer, *Generation of Pure-State Single-Photon Wavepackets by Conditional Preparation Based on Spontaneous Parametric Downconversion*, Laser Physics **15**, 146 (2005).

[80] T. Gerrits *et al.*, *Generation of degenerate, factorizable, pulsed squeezed light at telecom wavelengths*, Optics Express **19**, 24434 (2011).

[81] A. Eckstein, A. Christ, P. J. Mosley, and C. Silberhorn, *Highly Efficient Single-Pass Source of Pulsed Single-Mode Twin Beams of Light*, Physical Review Letters **106**, 013603 (2011).

[82] G. Harder, V. Ansari, B. Brecht, T. Dirmeier, C. Marquardt, and C. Silberhorn, *An optimized photon pair source for quantum circuits*, Optics Express **21**, 13975 (2013).

[83] S. Fasel, O. Alibart, S. Tanzilli, P. Baldi, A. Beveratos, N. Gisin, and H. Zbinden, *High-quality asynchronous heralded single-photon source at telecom wavelength*, New Journal of Physics **6**, 163 (2004).

[84] M. M. Weston *et al.*, *Efficient and pure femtosecond-pulse-length source of polarization-entangled photons*, Optics Express **24**, 10869 (2016).

[85] S. Ramelow *et al.*, *Highly efficient heralding of entangled single photons*, Optics Express **21**, 6707 (2013).

[86] F. Kaneda, K. Garay-Palmett, A. B. U'Ren, and P. G. Kwiat, *Heralded single-photon source utilizing highly nondegenerate, spectrally factorable spontaneous parametric down-conversion*, Optics Express **24**, 10733 (2016).

[87] R. J. A. Francis-Jones, R. A. Hoggarth, and P. J. Mosley, *All-fiber multiplexed source of high-purity single photons*, Optica **3**, 1270 (2016).

[88] J. B. Spring *et al.*, *Chip-based array of near-identical, pure, heralded single-photon sources*, Optica **4**, 90 (2017).

[89] P. S. Michelberger, M. Karpinski, I. A. Walmsley, and J. Nunn, *Optimal interfacing a GHz-bandwidth heralded single photon source with on-demand, broadband quantum memories*, `arXiv` 1709.03155v1 [quant-ph] (2017).

[90] A. McMillan, J. Fulconis, M. Halder, C. Xiong, J. Rarity, and W. Wadsworth, *Narrowband high-fidelity all-fibre source of heralded single photons at 1570 nm*, Optics Express **17**, 6156 (2009).

[91] A. Clark, B. Bell, J. Fulconis, M. M. Halder, B. Cemlyn, O. Alibart, C. Xiong, W. J. Wadsworth, and J. G. Rarity, *Intrinsically narrowband pair photon generation in microstructured fibres*, New Journal of Physics **13**, 065009 (2011).

[92] M. Rambach, A. Nikolova, T. J. Weinhold, and A. G. White, *Sub-megahertz linewidth single photon source*, APL Photonics **1**, 096101 (2016).

[93] B. Brecht, K.-H. Luo, H. Herrmann, and C. Silberhorn, *A versatile design for resonant guided-wave parametric down-conversion sources for quantum repeaters*, Applied Physics B **122**, 116 (2016).

[94] H.-S. Zhong *et al.*, *12-Photon Entanglement and Scalable Scattershot Boson Sampling with Optimal Entangled-Photon Pairs from Parametric Down-Conversion*, Physical Review Letters **121**, 250505 (2018).

[95] H. J. Kimble, M. Dagenais, and L. Mandel, *Photon Antibunching in Resonance Fluorescence*, Physical Review Letters **39**, 691 (1977).

[96] F. Diedrich and H. Walther, *Nonclassical radiation of a single stored ion*, Physical Review Letters **58**, 203 (1987).

[97] T. Basché, W. E. Moerner, M. Orrit, and H. Talon, *Photon antibunching in the fluorescence of a single dye molecule trapped in a solid*, Physical Review Letters **69**, 1516 (1992).

[98] P. Senellart, G. Solomon, and A. White, *High-performance semiconductor quantum-dot single-photon sources*, Nature Nanotechnology **12**, 1026 (2017), Review Article.

[99] P. Michler, A. Kiraz, C. Becher, W. V. Schoenfeld, P. M. Petroff, L. Zhang, E. Hu, and A. Imamoglu, *A Quantum Dot Single-Photon Turnstile Device*, Science **290**, 2282 (2000).

[100] D. Gammon, E. S. Snow, B. V. Shanabrook, D. S. Katzer, and D. Park, *Homogeneous Linewidths in the Optical Spectrum of a Single Gallium Arsenide Quantum Dot*, Science **273**, 87 (1996).

[101] M. Pelton, C. Santori, J. Vuc˘ković, B. Zhang, G. S. Solomon, J. Plant, and Y. Yamamoto, *Efficient Source of Single Photons: A Single Quantum Dot in a Micropost Microcavity*, Physical Review Letters **89**, 233602 (2002).

[102] G. S. Solomon, M. Pelton, and Y. Yamamoto, *Single-mode Spontaneous Emission from a Single Quantum Dot in a Three-Dimensional Microcavity*, Physical Review Letters **86**, 3903 (2001).

[103] A. Kiraz, P. Michler, C. Becher, B. Gayral, A. Imamolu, L. Zhang, E. Hu, W. V. Schoenfeld, and P. M. Petroff, *Cavity-quantum electrodynamics using a single InAs quantum dot in a microdisk structure*, Applied Physics Letters **78**, 3932 (2001).

[104] D. Englund, D. Fattal, E. Waks, G. Solomon, B. Zhang, T. Nakaoka, Y. Arakawa, Y. Yamamoto, and J. Vučković, *Controlling the Spontaneous Emission Rate of Single Quantum Dots in a Two-Dimensional Photonic Crystal*, Physical Review Letters **95**, 013904 (2005).

[105] X. Ding *et al.*, *On-Demand Single Photons with High Extraction Efficiency and Near-Unity Indistinguishability from a Resonantly Driven Quantum Dot in a Micropillar*, Physical Review Letters **116**, 020401 (2016).

[106] S. Unsleber, Y.-M. He, S. Gerhardt, S. Maier, C.-Y. Lu, J.-W. Pan, N. Gregersen, M. Kamp, C. Schneider, and S. Höfling, *Highly indistinguishable on-demand resonance fluorescence photons from a deterministic quantum dot micropillar device with 74% extraction efficiency*, Optics Express **24**, 8539 (2016).

[107] J. C. Loredo *et al.*, *Scalable performance in solid-state single-photon sources*, Optica **3**, 433 (2016).

[108] E. Schöll *et al.*, *Resonance Fluorescence of GaAs Quantum Dots with Near-Unity Photon Indistinguishability*, Nano Letters **19**, 2404 (2019).

[109] W. A. Gambling, *The rise and rise of optical fibers*, IEEE Journal of Selected Topics in Quantum Electronics **6**, 1084 (2000).

[110] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication*, Physical Review Letters **81**, 5932 (1998).

[111] M. Żukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert, *"Event-ready-detectors" Bell experiment via entanglement swapping*, Physical Review Letters **71**, 4287 (1993).

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

[113] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, *Long-distance quantum communication with atomic ensembles and linear optics*, Nature **414**, 413 (2001).

[114] A. Imamoğlu, *High Efficiency Photon Counting Using Stored Light*, Physical Review Letters **89**, 163602 (2002).

[115] D. F. V. James and P. G. Kwiat, *Atomic-Vapor-Based High Efficiency Optical Detectors with Photon Number Resolution*, Physical Review Letters **89**, 183601 (2002).

[116] D. Budker and M. Romalis, *Optical magnetometry*, Nature Physics **3**, 227 (2007), Review Article.

[117] J. Appel, P. J. Windpassinger, D. Oblak, U. B. Hoff, N. Kjærgaard, and E. S. Polzik, *Mesoscopic atomic entanglement for precision measurements beyond the standard quantum limit*, Proceedings of the National Academy of Sciences **106**, 10960 (2009).

[118] 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*, Physical Review Letters **110**, 133601 (2013).

[119] F. Bussires, N. Sangouard, M. Afzelius, H. de Riedmatten, C. Simon, and W. Tittel, *Prospective applications of optical quantum memories*, Journal of Modern Optics **60**, 1519 (2013).

[120] S. Massar and S. Popescu, *Optimal Extraction of Information from Finite Quantum Ensembles*, Physical Review Letters **74**, 1259 (1995).

[121] P. W. Shor, *Scheme for reducing decoherence in quantum computer memory*, Physical Review A **52**, R2493 (1995).

[122] M. Gündoğan, P. M. Ledingham, K. Kutluer, M. Mazzera, and H. de Riedmatten, *Solid State Spin-Wave Quantum Memory for Time-Bin Qubits*, Physical Review Letters **114**, 230501 (2015).

[123] K. T. Kaczmarek *et al.*, *High-speed noise-free optical quantum memory*, Physical Review A **97**, 042316 (2018).

[124] C. Simon *et al.*, *Quantum memories*, The European Physical Journal D **58**, 1 (2010).

[125] W. Munro, K. A. Harrison, A. M. Stephens, S. J. Devitt, and K. Nemoto, *From quantum multiplexing to high-performance quantum networking*, Nature Photonics **4**, 792 (2010).

[126] F. M. Miatto, M. Epping, and N. Lütkenhaus, *Hamiltonians for one-way quantum repeaters*, Quantum **2**, 75 (2018).

[127] D. V. Vasilyev, I. V. Sokolov, and E. S. Polzik, *Quantum memory for images: A quantum hologram*, Physical Review A **77**, 020302 (2008).

[128] A. Grodecka-Grad, E. Zeuthen, and A. S. Sørensen, *High-Capacity Spatial Multimode Quantum Memories Based on Atomic Ensembles*, Physical Review Letters **109**, 133601 (2012).

[129] W. Tittel, M. Afzelius, T. Chanelire, R. Cone, S. Krll, S. Moiseev, and M. Sellars, *Photon-echo quantum memory in solid state systems*, Laser & Photonics Reviews **4**, 244 (2010).

[130] N. Sinclair *et al.*, *Spectral Multiplexing for Scalable Quantum Photonics using an Atomic Frequency Comb Quantum Memory and Feed-Forward Control*, Physical Review Letters **113**, 053603 (2014).

[131] C. Simon, H. de Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, *Quantum Repeaters with Photon Pair Sources and Multimode Memories*, Physical Review Letters **98**, 190503 (2007).

[132] M. Gündogan, P. M. Ledingham, A. Almasi, M. Cristiani, and H. de Riedmatten, *Quantum Storage of a Photonic Polarization Qubit in a Solid*, Physical Review Letters **108**, 190504 (2012).

[133] Z. Zheng, O. Mishina, N. Treps, and C. Fabre, *Atomic quantum memory for multimode frequency combs*, Physical Review A **91**, 031802 (2015).

[134] S. Gao, O. Lazo-Arjona, B. Brecht, K. T. Kaczmarek, S. E. Thomas, J. Nunn, P. M. Ledingham, D. J. Saunders, and I. A. Walmsley, *Optimal Coherent Filtering for Single Noisy Photons*, arXiv 1902.07720 [quant-ph] (2019).

[135] K. Heshami, D. G. England, P. C. Humphreys, P. J. Bustard, V. M. Acosta, J. Nunn, and B. J. Sussman, *Quantum memories: emerging applications and recent advances*, Journal of Modern Optics **63**, 2005 (2016).

[136] T. B. Pittman and J. D. Franson, *Cyclical quantum memory for photonic qubits*, Physical Review A **66**, 062302 (2002).

[137] X. Li, P. L. Voss, J. Chen, J. E. Sharping, and P. Kumar, *Storage and long-distance distribution of telecommunications-band polarization entanglement generated in an optical fiber*, Optics Letters **30**, 1201 (2005).

[138] F. Kaneda, F. Xu, J. Chapman, and P. G. Kwiat, *Quantum-memory-assisted multi-photon generation for efficient quantum information processing*, Optica **4**, 1034 (2017).

[139] C. Langer *et al.*, *Long-Lived Qubit Memory Using Atomic Ions*, Physical Review Letters **95**, 060502 (2005).

[140] H. P. Specht, C. Nölleke, A. Reiserer, M. Uphoff, E. Figueroa, S. Ritter, and G. Rempe, *A single-atom quantum memory*, Nature **473**, 190 (2011).

[141] G. D. Fuchs, G. Burkard, P. V. Klimov, and D. D. Awschalom, *A quantum memory intrinsic to single nitrogen-vacancy centres in diamond*, Nature Physics **7**, 789 (2011).

[142] D. D. Sukachev, A. Sipahigil, C. T. Nguyen, M. K. Bhaskar, R. E. Evans, F. Jelezko, and M. D. Lukin, *Silicon-Vacancy Spin Qubit in Diamond: A Quantum Memory Exceeding 10 ms with Single-Shot State Readout*, Physical Review Letters **119**, 223602 (2017).

[143] A. Sipahigil *et al.*, *An integrated diamond nanophotonics platform for quantum-optical networks*, Science **354**, 847 (2016).

[144] Y. Wang, J. Li, S. Zhang, K. Su, Y. Zhou, K. Liao, S. Du, H. Yan, and S.-L. Zhu, *Efficient quantum memory for single-photon polarization qubits*, Nature Photonics **13**, 346 (2019).

[145] J. J. Longdell, E. Fraval, M. J. Sellars, and N. B. Manson, *Stopped Light with Storage Times Greater than One Second Using Electromagnetically Induced Transparency in a Solid*, Physical Review Letters **95**, 063601 (2005).

[146] D. G. England, K. A. G. Fisher, J.-P. W. MacLean, P. J. Bustard, R. Lausten, K. J. Resch, and B. J. Sussman, *Storage and Retrieval of THz-Bandwidth Single Photons Using a Room-Temperature Diamond Quantum Memory*, Physical Review Letters **114**, 053602 (2015).

[147] I. D. Abella, N. A. Kurnit, and S. R. Hartmann, *Photon Echoes*, Physical Review **141**, 391 (1966).

[148] G. Hétet, J. J. Longdell, A. L. Alexander, P. K. Lam, and M. J. Sellars, *Electro-Optic Quantum Memory for Light Using Two-Level Atoms*, Physical Review Letters **100**, 023601 (2008).

[149] S. A. Moiseev and S. Kröll, *Complete Reconstruction of the Quantum State of a Single-Photon Wave Packet Absorbed by a Doppler-Broadened Transition*, Physical Review Letters **87**, 173601 (2001).

[150] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler, *Unconditional room-temperature quantum memory*, Nature Physics **7**, 794 (2011).

[151] B. M. Sparkes, J. Bernu, M. Hosseini, J. Geng, Q. Glorieux, P. A. Altin, P. K. Lam, N. P. Robins, and B. C. Buchler, *Gradient echo memory in an ultra-high optical depth cold atomic ensemble*, New Journal of Physics **15**, 085027 (2013).

[152] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin, *Multimode quantum memory based on atomic frequency combs*, Physical Review A **79**, 052329 (2009).

[153] M. Sabooni, Q. Li, S. Kröll, and L. Rippe, *Efficient Quantum Memory Using a Weakly Absorbing Sample*, Physical Review Letters **110**, 133604 (2013).

[154] M. Afzelius *et al.*, *Demonstration of Atomic Frequency Comb Memory for Light with Spin-Wave Storage*, Physical Review Letters **104**, 040503 (2010).

[155] P. Jobez, I. Usmani, N. Timoney, C. Laplane, N. Gisin, and M. Afzelius, *Cavity-enhanced storage in an optical spin-wave memory*, New Journal of Physics **16**, 083005 (2014).

[156] C. Laplane, P. Jobez, J. Etesse, N. Timoney, N. Gisin, and M. Afzelius, *Multiplexed on-demand storage of polarization qubits in a crystal*, New Journal of Physics **18**, 013006 (2015).

[157] M. Gündogan, P. M. Ledingham, K. Kutluer, M. Mazzera, and H. de Riedmatten, *Solid State Spin-Wave Quantum Memory for Time-Bin Qubits*, Physical Review Letters **114**, 230501 (2015).

[158] A. Seri, A. Lenhard, D. Rieländer, M. Gündoğan, P. M. Ledingham, M. Mazzera, and H. de Riedmatten, *Quantum Correlations between Single Telecom Photons and a Multimode On-Demand Solid-State Quantum Memory*, Physical Review X **7**, 021028 (2017).

[159] I. Usmani, M. Afzelius, H. de Riedmatten, and N. Gisin, *Mapping multiple photonic qubits into and out of one solid-state atomic ensemble*, Nature Communications **1**, 12 (2010).

[160] M. Bonarota, J.-L. L. Gout, and T. Chanelière, *Highly multimode storage in a crystal*, New Journal of Physics **13**, 013013 (2011).

[161] E. Saglamyurek, N. Sinclair, J. Jin, J. A. Slater, D. Oblak, F. Bussières, M. George, R. Ricken, W. Sohler, and W. Tittel, *Broadband waveguide quantum memory for entangled photons*, Nature **469**, 512 (2011).

[162] N. Sinclair *et al.*, *Spectral Multiplexing for Scalable Quantum Photonics using an Atomic Frequency Comb Quantum Memory and Feed-Forward Control*, Physical Review Letters **113**, 053603 (2014).

[163] K.-J. Boller, A. Imamoğlu, and S. E. Harris, *Observation of electromagnetically induced transparency*, Physical Review Letters **66**, 2593 (1991).

[164] Dutton, Zachary, Ginsberg, Naomi S., Slowe, Christopher, and Hau, Lene Vestergaard, *The art of taming light: ultra-slow and stopped light*, Europhysics News **35**, 33 (2004).

[165] C. Liu, Z. Dutton, C. H. Behroozi, and L. V. Hau, *Observation of coherent optical information storage in an atomic medium using halted light pulses*, Nature **409**, 490 (2001).

[166] M. Fleischhauer and M. D. Lukin, *Dark-State Polaritons in Electromagnetically Induced Transparency*, Physical Review Letters **84**, 5094 (2000).

[167] A. V. Gorshkov, A. André, M. Fleischhauer, A. S. Sørensen, and M. D. Lukin, *Universal Approach to Optimal Photon Storage in Atomic Media*, Physical Review Letters **98**, 123601 (2007).

[168] D. F. Phillips, A. Fleischhauer, A. Mair, R. L. Walsworth, and M. D. Lukin, *Storage of Light in Atomic Vapor*, Physical Review Letters **86**, 783 (2001).

[169] A. V. Turukhin, V. S. Sudarshanam, M. S. Shahriar, J. A. Musser, B. S. Ham, and P. R. Hemmer, *Observation of Ultraslow and Stored Light Pulses in a Solid*, Physical Review Letters **88**, 023602 (2001).

[170] G. Heinze, C. Hubrich, and T. Halfmann, *Stopped Light and Image Storage by Electromagnetically Induced Transparency up to the Regime of One Minute*, Physical Review Letters **111**, 033601 (2013).

[171] T. Chanelière, D. N. Matsukevich, S. D. Jenkins, S.-Y. Lan, T. A. B. Kennedy, and A. Kuzmich, *Storage and retrieval of single photons transmitted between remote quantum memories*, Nature **438**, 833 (2005).

[172] M. D. Eisaman, A. André, F. Massou, M. Fleischhauer, A. S. Zibrov, and M. D. Lukin, *Electromagnetically induced transparency with tunable single-photon pulses*, Nature **438**, 837 (2005).

[173] E. Distante, P. Farrera, A. Padrón-Brito, D. Paredes-Barato, G. Heinze, and H. de Riedmatten, *Storing single photons emitted by a quantum memory on a highly excited Rydberg state*, Nature Communications **8**, 14072 (2017).

[174] E. Saglamyurek, T. Hrushevskyi, A. Rastogi, K. Heshami, and L. J. LeBlanc, *Coherent storage and manipulation of broadband photons via dynamically controlled Autler-Townes splitting*, Nature Photonics **12**, 774 (2018).

[175] A. E. Kozhekin, K. Mølmer, and E. Polzik, *Quantum memory for light*, Physical Review A **62**, 033809 (2000).

[176] J. Nunn, I. A. Walmsley, M. G. Raymer, K. Surmacz, F. C. Waldermann, Z. Wang, and D. Jaksch, *Mapping broadband single-photon wave packets into an atomic memory*, Physical Review A **75**, 011401 (2007).

[177] K. F. Reim, J. Nunn, V. O. Lorenz, B. J. Sussman, K. C. Lee, N. K. Langford, D. Jaksch, and I. A. Walmsley, *Towards high-speed optical quantum memories*, Nature Photonics **4**, 218 (2010).

[178] P. J. Bustard, D. G. England, K. Heshami, C. Kupchak, and B. J. Sussman, *Reducing noise in a Raman quantum memory*, Optics Letters **41**, 5055 (2016).

[179] J. Guo, X. Feng, P. Yang, Z. Yu, L. Q. Chen, C.-H. Yuan, and W. Zhang, *High-performance Raman quantum memory with optimal control in room temperature atoms*, Nature Communications **10**, 148 (2019).

[180] K. F. Reim, P. Michelberger, K. C. Lee, J. Nunn, N. K. Langford, and I. A. Walmsley, *Single-Photon-Level Quantum Memory at Room Temperature*, Physical Review Letters **107**, 053603 (2011).

[181] D. G. England *et al.*, *High-fidelity polarization storage in a gigahertz bandwidth quantum memory*, Journal of Physics B: Atomic, Molecular and Optical Physics **45**, 124008 (2012).

[182] D.-S. Ding, W. Zhang, Z.-Y. Zhou, S. Shi, B.-S. Shi, and G.-C. Guo, *Raman quantum memory of photonic polarized entanglement*, Nature Photonics **9**, 332 (2015).

[183] P. S. Michelberger *et al.*, *Interfacing GHz-bandwidth heralded single photons with a warm vapour Raman memory*, New Journal of Physics **17**, 043006 (2015).

[184] R. Finkelstein, E. Poem, O. Michel, O. Lahad, and O. Firstenberg, *Fast, noise-free memory for photon synchronization at room temperature*, Science Advances **4**, 1 (2018).

[185] S.-J. Yang, X.-J. Wang, X.-H. Bao, and J.-W. Pan, *An efficient quantum light-matter interface with sub-second lifetime*, Nature Photonics **10**, 381 (2016).

[186] R. Chrapkiewicz, M. Dabrowski, and W. Wasilewski, *High-Capacity Angularly Multiplexed Holographic Memory Operating at the Single-Photon Level*, Physical Review Letters **118**, 063603 (2017).

[187] Y.-F. Pu, N. Jiang, W. Chang, H.-X. Yang, C. Li, and L.-M. Duan, *Experimental realization of a multiplexed quantum memory with 225 individually accessible memory cells*, Nature Communications **8**, 15359 (2017).

[188] N. V. Corzo, J. Raskop, A. Chandra, A. S. Sheremet, B. Gouraud, and J. Laurat, *Waveguide-coupled single collective excitation of atomic arrays*, Nature **566**, 359 (2019).

[189] J. Nunn, *Quantum Memory in Atomic Ensembles*, PhD thesis, University of Oxford, 2008.

[190] J. Munns, *Optimisation and applications of a Raman quantum memory for temporal modes of light*, PhD thesis, Imperial College London, 2018.

[191] K. Hammerer, A. S. Sørensen, and E. S. Polzik, *Quantum interface between light and atomic ensembles*, Reviews of Modern Physics **82**, 1041 (2010).

[192] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. II. Free-space model*, Physical Review A **76**, 033805 (2007).

[193] J. Nunn, K. Reim, K. C. Lee, V. O. Lorenz, B. J. Sussman, I. A. Walmsley, and D. Jaksch, *Multimode Memories in Atomic Ensembles*, Physical Review Letters **101**, 260502 (2008).

[194] P. Walther, M. D. Eisaman, A. Andre, F. Massou, M. Fleischhauer, A. S. Zibrov, and M. D. Lukin, *Generation of narrow-bandwidth single photons using electromagnetically induced transparency in atomic ensembles*, International Journal of Quantum Information **5**, 51 (2007).

[195] E. Poem, C. Weinzetl, J. Klatzow, K. T. Kaczmarek, J. H. D. Munns, T. F. M. Champion, D. J. Saunders, J. Nunn, and I. A. Walmsley, *Broadband noise-free optical quantum memory with neutral nitrogen-vacancy centers in diamond*, Physical Review B **91**, 205108 (2015).

[196] I. Vurgaftman and M. Bashkansky, *Suppressing four-wave mixing in warm-atomic-vapor quantum memory*, Physical Review A **87**, 063836 (2013).

[197] S. E. Thomas, High efficiency broadband raman memory in warm caesium vapour using quenching, Master's thesis, Imperial College London, 2015.

[198] S. E. Thomas, J. H. D. Munns, K. T. Kaczmarek, C. Qiu, B. Brecht, A. Feizpour, P. M. Ledingham, I. A. Walmsley, J. Nunn, and D. J. Saunders, *High efficiency Raman memory by suppressing radiation trapping*, New Journal of Physics **19**, 063034 (2017).

[199] D. A. Steck, *Cesium D Line Data*, http://steck.us/alkalidata/cesiumnumbers.1.6.pdf.

[200] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, *Laser phase and frequency stabilization using an optical resonator*, Applied Physics B **31**, 97 (1983).

[201] W. Happer, *Optical Pumping*, Review of Modern Physics **44**, 169 (1972).

[202] K. T. Kaczmarek, D. J. Saunders, M. R. Sprague, W. S. Kolthammer, A. Feizpour, P. M. Ledingham, B. Brecht, E. Poem, I. A. Walmsley, and J. Nunn, *Ultrahigh and persistent optical depths of cesium in Kagomé-type hollow-core photonic crystal fibers*, Optics Letters **40**, 5582 (2015).

[203] M. A. Rosenberry, J. P. Reyes, D. Tupa, and T. J. Gay, *Radiation trapping in rubidium optical pumping at low buffer-gas pressures*, Physical Review A **75**, 023401 (2007).

[204] G. A. Pitz, *Collisional dynamics of the cesium D_1 and D_2 transitions*, PhD thesis, Air Force Institute of Technology, 2010.

[205] J. Nunn *et al.*, *Theory of noise suppression in Λ-type quantum memories by means of a cavity*, Physical Review A **96**, 012338 (2017).

[206] C. F. McCormick, A. M. Marino, V. Boyer, and P. D. Lett, *Strong low-frequency quantum correlations from a four-wave-mixing amplifier*, Physical Review A **78**, 043816 (2008).

[207] D. J. Saunders, J. H. D. Munns, T. F. M. Champion, C. Qiu, K. T. Kaczmarek, E. Poem, P. M. Ledingham, I. A. Walmsley, and J. Nunn, *Cavity-Enhanced Room-Temperature Broadband Raman Memory*, Physical Review Letters **116**, 090501 (2016).

[208] T. F. M. Champion, *Towards storage and retrieval of non-classical light in a broadband quantum memory - An investigation of free-space and cavity Raman memories*, PhD thesis, University of Oxford, 2015.

[209] M. D. Eisaman, J. Fan, A. Migdall, and S. V. Polyakov, *Invited Review Article: Single-photon sources and detectors*, Review of Scientific Instruments **82**, 071101 (2011).

[210] W. Boyd, *Nonlinear Optics* (Academic Press, 2003).

[211] G. Romanov, C. OBrien, and I. Novikova, *Suppression of the four-wave mixing amplification via Raman absorption*, Journal of Modern Optics **63**, 2048 (2016).

[212] N. Prajapati, G. Romanov, and I. Novikova, *Suppression of four-wave mixing in hot rubidium vapor using ladder scheme Raman absorption*, Journal of the Optical Society of America B **34**, 1994 (2017).

[213] V. Popkov, S. Essink, C. Presilla, and G. Schütz, *Effective quantum Zeno dynamics in dissipative quantum systems*, Physical Review A **98**, 052110 (2018).

[214] B. Brecht, D. V. Reddy, C. Silberhorn, and M. G. Raymer, *Photon Temporal Modes: A Complete Framework for Quantum Information Science*, Physical Review X **5**, 041017 (2015).

[215] C. Rosales-Guzmn, B. Ndagano, and A. Forbes, *A review of complex vector light fields and their applications*, Journal of Optics **20**, 123001 (2018).

[216] N. T. Islam, C. C. W. Lim, C. Cahall, J. Kim, and D. J. Gauthier, *Provably secure and high-rate quantum key distribution with time-bin qudits*, Science Advances **3**, 11 (2017).

[217] A. Sit *et al.*, *High-dimensional intracity quantum cryptography with structured photons*, Optica **4**, 1006 (2017).

[218] Y. Ren *et al.*, *Spatially multiplexed orbital-angular-momentum-encoded single photon and classical channels in a free-space optical communication link*, Optics Letters **42**, 4881 (2017).

[219] V. Ansari, J. M. Donohue, B. Brecht, and C. Silberhorn, *Tailoring nonlinear processes for quantum optics with pulsed temporal-mode encodings*, Optica **5**, 534 (2018).

[220] J. A. Salehi, *Code division multiple-access techniques in optical fiber networks. I. Fundamental principles*, IEEE Transactions on Communications **37**, 824 (1989).

[221] I. A. Walmsley and V. Wong, *Characterization of the electric field of ultrashort optical pulses*, Journal of the Optical Society of America B **13**, 2453 (1996).

[222] B. Srivathsan, G. K. Gulati, A. Cerè, B. Chng, and C. Kurtsiefer, *Reversing the Temporal Envelope of a Heralded Single Photon using a Cavity*, Physical Review Letters **113**, 163601 (2014).

[223] I. Novikova, N. B. Phillips, and A. V. Gorshkov, *Optimal light storage with full pulse-shape control*, Physical Review A **78**, 021802 (2008).

[224] M. G. Raymer, S. J. van Enk, C. J. McKinstrie, and H. J. McGuinness, *Interference of two photons of different color*, Optics Communications **283**, 747 (2010).

[225] D. V. Reddy, M. G. Raymer, and C. J. McKinstrie, *Sorting photon wave packets using temporal-mode interferometry based on multiple-stage quantum frequency conversion*, Physical Review A **91**, 012323 (2015).

[226] M. Allgaier, V. Ansari, L. Sansoni, C. Eigner, V. Quiring, R. Ricken, G. Harder, B. Brecht, and C. Silberhorn, *Highly efficient frequency conversion with bandwidth compression of quantum light*, Nature Communications **8**, 14288 (2017).

[227] V. Ansari, G. Harder, M. Allgaier, B. Brecht, and C. Silberhorn, *Temporal-mode measurement tomography of a quantum pulse gate*, Physical Review A **96**, 063817 (2017).

[228] A. Christ, B. Brecht, W. Mauerer, and C. Silberhorn, *Theory of quantum frequency conversion and type-II parametric down-conversion in the high-gain regime*, New Journal of Physics **15**, 053038 (2013).

[229] N. Quesada and J. E. Sipe, *High efficiency in mode-selective frequency conversion*, Optics Letters **41**, 364 (2016).

[230] F. Krumm, J. Sperling, and W. Vogel, *Multitime correlation functions in nonclassical stochastic processes*, Physical Review A **93**, 063843 (2016).

[231] D. V. Reddy and M. G. Raymer, *Engineering temporal-mode-selective frequency conversion in nonlinear optical waveguides: from theory to experiment*, Optics Express **25**, 12952 (2017).

[232] K. A. G. Fisher, D. G. England, J.-P. W. MacLean, P. J. Bustard, K. J. Resch, and B. J. Sussman, *Frequency and bandwidth conversion of single photons in a room-temperature diamond quantum memory*, Nature Communications **7**, 11200 (2016).

[233] C. K. Law, I. A. Walmsley, and J. H. Eberly, *Continuous Frequency Entanglement: Effective Finite Hilbert Space and Entropy Control*, Physical Review Letters **84**, 5304 (2000).

[234] D. G. England, K. A. G. Fisher, J.-P. W. MacLean, P. J. Bustard, K. Heshami, K. J. Resch, and B. J. Sussman, *Phonon-Mediated Nonclassical Interference in Diamond*, Physical Review Letters **117**, 073603 (2016).

[235] C. Iaconis and I. A. Walmsley, *Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses*, Optics Letters **23**, 792 (1998).

[236] E. A. Goldschmidt, F. Piacentini, I. R. Berchera, S. V. Polyakov, S. Peters, S. Kück, G. Brida, I. P. Degiovanni, A. Migdall, and M. Genovese, *Mode reconstruction of a light field by multiphoton statistics*, Physical Review A **88**, 013822 (2013).

Appendix A:

Green's Function analysis

To experimentally quantify the amount of noise generated by the quantum memory, we use two figures of merit. Firstly we consider the μ_1 parameter which is defined as $\mu_1 = N_{\text{noise}}/\eta_{\text{mem}}$, where N_{noise} is the average number of noise photons per memory interaction. This parameter represents the number of photons you need at the input to have an signal-to-noise ratio (SNR) of 1 on the output, and is a useful parameter to quantify the total amount of noise generated by the memory and make comparisons between different protocols.

Secondly, for a more detailed analysis of the noise processes and how these affect the photon number statistics of the stored state, we consider the second-order autocorrelation of the output photonic state. This is defined as

$$g^{(2)}_{\text{out}} = \frac{\iint \mathrm{d}t\,\mathrm{d}t' \left\langle S^{\dagger}_{\text{out}}(t')S^{\dagger}_{\text{out}}(t)S_{\text{out}}(t)S_{\text{out}}(t')\right\rangle}{\left(\int \mathrm{d}t \left\langle S^{\dagger}_{\text{out}}(t)S_{\text{out}}(t)\right\rangle\right)^2}, \tag{A.1}$$

and $g^{(2)} < 1$ signifies a field with non-classical statistics.

There are several noise processes that can contribute to the photon number statistics. Firstly we consider noise arising from spontaneous Raman scattering processes, N_{SRS}, which can be due to four-wave mixing noise, or spontaneous Stokes scattering due to the initial unpumped proportion of the population in $|3\rangle$, α. These noise processes are accounted for by the equations of motion, and therefore we can derive an expression for their contribution to $g^{(2)}_{\text{out}}$ via Green's function analysis of Equations 5.1-5.3. Noise sources that are not included in Equations 5.1-5.3 are those arising from the excited state, such as broadband collision-induced fluorescence, N_{F}, that is not sufficiently filtered from detection and direct leakage of the strong control field due to insufficient filtering, N_{L}. We note that, in our experimental set-up, the filtering of the control field is sufficiently strong that $N_{\text{L}} = 0$, which is confirmed by removing the memory vapour cell and measuring that the noise level when there are no atoms present is within the dark

counts of the detectors. We mix the output statistics of the fluorescence noise process with those derived from the equations of motion to get an expression for the statistics of the total retrieved state [236].

An expression for the $g^{(2)}$ of the retrieved state can be derived via Green's function analysis of Equations 5.1-5.3. The results presented below have been derived by T.M.H. and O.L.A.

The equations of motion (Equations 5.1-5.3) are linear so the resulting evolution may be characterised by Green's function mappings G_{ij} from initial mode i to final mode j as:

$$\begin{aligned}
S_{\text{out}}(t) &= \int_{-\infty}^{\infty} \mathrm{d}t_1\, G_{ss}(t,t_1)S_{\text{in}}(t_1) + \int_{-\infty}^{\infty} \mathrm{d}t_1\, G_{sa^\dagger}(t,t_1)A^\dagger_{\text{in}}(t_1) + \int_0^1 \mathrm{d}z'\, G_{sb}(t,z_1)B_{\text{in}}(z_1) \\
A^\dagger_{\text{out}}(t) &= \int_{-\infty}^{\infty} \mathrm{d}t_1\, G_{a^\dagger s}(t,t_1)S_{\text{in}}(t_1) + \int_{-\infty}^{\infty} \mathrm{d}t_1\, G_{a^\dagger a^\dagger}(t,t_1)A^\dagger_{\text{in}}(t_1) + \int_0^1 \mathrm{d}z_1\, G_{a^\dagger b}(t,z_1)B_{\text{in}}(z_1) \\
B_{\text{out}}(z) &= \int_{-\infty}^{\infty} \mathrm{d}t_1\, G_{bs}(z,t_1)S_{\text{in}}(t_1) + \int_{-\infty}^{\infty} \mathrm{d}t_1\, G_{ba^\dagger}(z,t_1)A^\dagger_{\text{in}}(t_1) + \int_0^1 \mathrm{d}z_1\, G_{bb}(z,z_1)B_{\text{in}}(z_1).
\end{aligned}$$

To simplify notation we define $G_{sati} \equiv G_{sa^\dagger}(t,t_i)$ and $S_i \equiv S_{\text{in}}(t_i)$, and take out the explicit z- or t-dependence in the integrals by letting $\mathrm{d}_1 \equiv \mathrm{d}t_1$ or $\mathrm{d}z_1$ as appropriate. This simplifies the first line above to:

$$S_{\text{out}}(t) = \int \mathrm{d}_1 \left(G_{sst1}S_1 + G_{sa^\dagger t1}A^\dagger_1 + G_{sbt1}B_1 \right). \tag{A.2}$$

To arrive at an expression for the second order auto-correlation we want to consider

$$\begin{aligned}
S^\dagger_{t'}S^\dagger_t S_t S_{t'} &= \int \left(G^*_{sst'1}S^\dagger_1 + G^*_{sa^\dagger t'1}A_1 + G^*_{sbt'1}B^\dagger_1 \right) \mathrm{d}_1 \times \int \left(G^*_{sst2}S^\dagger_2 + G^*_{sa^\dagger t2}A_2 + G^*_{sbt2}B^\dagger_2 \right) \mathrm{d}_2 \\
&\times \int \left(G_{sst3}S_3 + G_{sa^\dagger t3}A^\dagger_3 + G_{sbt3}B_3 \right) \mathrm{d}_3 \times \int \left(G_{sst'4}S_4 + G_{sa^\dagger t'4}A^\dagger_4 + G_{sbt'4}B_4 \right) \mathrm{d}_4,
\end{aligned} \tag{A.3}$$

and then finally integrate this expression over the variables t and t'.

We consider the input signal to be a superposition of Fock states:

$$\begin{aligned}
|s_{\text{in}}\rangle &= \sum_n c_n |n\rangle, \\
N_{\text{in}} &\equiv \int \mathrm{d}t\, \langle s_{\text{in}}| S^\dagger_t S_t |s_{\text{in}}\rangle.
\end{aligned}$$

We take the input anti-Stokes field to be vacuum, $|a_{\text{in}}\rangle = |0_a\rangle$, and the initial spin wave to have population $\langle B^\dagger_{\text{in}} B_{\text{in}}\rangle = \alpha$, given by number of unpumped atoms in initial state $|3\rangle$. The full input state is described as $|\Psi_{\text{in}}\rangle = |s_{\text{in}}\rangle |0_a\rangle |b_{\text{in}}\rangle$ and the number of photons retrieved from the memory is given by

$$
\begin{aligned}
N_{\text{out}} &= \int \mathrm{d}t\ \langle \Psi_{\text{in}}|\, S^{\dagger}_{\text{out}}(t) S_{\text{out}}(t)\, |\Psi_{\text{in}}\rangle \\
&= \int \mathrm{d}t\ \left\langle \Psi_{\text{in}} \left| \int \mathrm{d}_1 \left(G^{*}_{sst1} S^{\dagger}_1 + G^{*}_{sa^{\dagger}t1} A_1 + G^{*}_{sbt1} B^{\dagger}_1 \right) \times \int \mathrm{d}_2 \left(G_{sst2} S_2 + G_{sa^{\dagger}t2} A^{\dagger}_2 + G_{sbt2} B_2 \right) \right| \Psi_{\text{in}} \right\rangle \\
&= \int \mathrm{d}t \int \mathrm{d}_{12} \left(G^{*}_{sst1} G_{sst2} \langle s_{\text{in}}|\, S^{\dagger}_1 S_2\, |s_{\text{in}}\rangle + G^{*}_{sa^{\dagger}t1} G_{sa^{\dagger}t2} \langle 0_a|\, A_1 A^{\dagger}_2\, |0_a\rangle + G^{*}_{sbt1} G_{sbt2} \langle b_{\text{in}}|\, B^{\dagger}_1 B_2\, |b_{\text{in}}\rangle \right) \\
&\equiv N_{\text{mem}} + N^{AS}_{\text{spont}} + N^{P}_{\text{spont}}
\end{aligned}
\tag{A.4}
$$

where the first term describes output photons due to the desired memory interaction, the second term describes output photons arising from spontaneous four-wave mixing, and the third term is spontaneous Raman scattering due to the initial occupation of the spin wave from imperfect optical pumping.

The second order autocorrelation function of the input signal is:

$$
g^{(2)}_{\text{in}} = \frac{\iint \mathrm{d}t\, \mathrm{d}t'\ \langle \psi_{\text{in}}|\, S^{\dagger}_{t'} S^{\dagger}_t S_t S_{t'}\, |\psi_{\text{in}}\rangle}{\left(\int \mathrm{d}t\ \langle \psi_{\text{in}}|\, S^{\dagger}_t S_t\, |\psi_{\text{in}}\rangle \right)^2} = \frac{\iint \mathrm{d}t\, \mathrm{d}t'\ \langle \psi_{\text{in}}|\, S^{\dagger}_{t'} S^{\dagger}_t S_t S_{t'}\, |\psi_{\text{in}}\rangle}{N^2_{\text{in}}},
$$

and so

$$
\iint \mathrm{d}t\, \mathrm{d}t'\ \langle \psi_{\text{in}}|\, S^{\dagger}_{t'} S^{\dagger}_t S_t S_{t'}\, |\psi_{\text{in}}\rangle = N^2_{\text{in}} g^{(2)}_{\text{in}}.
$$

We now consider the non-zero terms in Equation A.3. First, the diagonal terms are given by:

$$
\begin{aligned}
S: \quad & \int \mathrm{d}t \int \mathrm{d}_{1234}\, G^{*}_{sst1} G^{*}_{sst2} G_{sst3} G_{sst4} \langle s_{in}|\, S^{\dagger}_1 S^{\dagger}_2 S_3 S_4\, |s_{in}\rangle \\
&= \int \mathrm{d}t \int \mathrm{d}_{1234}\, G^{*}_{sst1} G^{*}_{sst2} G_{sst3} G_{sst4} \langle s_{in}|\, S^{\dagger}_1 S^{\dagger}_2 S_2 S_1\, |s_{in}\rangle\, \delta_{14}\delta_{23} \\
&= N^2_{\text{in}} g^{(2)}_{\text{in}} \int \mathrm{d}t \int \mathrm{d}_{12}\, G^{*}_{sst1} G^{*}_{sst2} G_{sst2} G_{sst1} \\
&= N^2_{\text{in}} g^{(2)}_{\text{in}} \int \mathrm{d}t \int \mathrm{d}_1\ |G_{sst1} G_{sst1}|^2 \\
&\equiv N^2_{\text{in}} g^{(2)}_{\text{in}} \mathcal{G}_{ss},
\end{aligned}
$$

$$
\begin{aligned}
A: \quad & \int \mathrm{d}t \int \mathrm{d}_{1234}\, G^{*}_{sa^{\dagger}t1} G^{*}_{sa^{\dagger}t2} G_{sa^{\dagger}t3} G_{sa^{\dagger}t4} \langle 0_a|\, A_1 A_2 A^{\dagger}_3 A^{\dagger}_4\, |0_a\rangle \\
&= \int \mathrm{d}t \int \mathrm{d}_{1234}\, G^{*}_{sa^{\dagger}t1} G^{*}_{sa^{\dagger}t2} G_{sa^{\dagger}t3} G_{sa^{\dagger}t4} (\delta_{13}\delta_{24} + \delta_{14}\delta_{23}) \\
&= 2 \int \mathrm{d}t \int \mathrm{d}_1\ |G_{sa^{\dagger}t1} G_{sa^{\dagger}t1}|^2 \\
&\equiv 2\mathcal{G}_{sa^{\dagger}},
\end{aligned}
$$

$$
\begin{aligned}
B: \quad & \int \mathrm{d}t \int \mathrm{d}_{1234}\, G^*_{sbt1} G^*_{sbt2} G_{sbt3} G_{sbt4} \left\langle b_{in}\right| B_1^\dagger B_2^\dagger B_3 B_4 \left|b_{in}\right\rangle \\
& = \int \mathrm{d}t \int \mathrm{d}_{1234}\, G^*_{sbt1} G^*_{sbt3} G_{sbt2} G_{sbt4} \left\langle b_{in}\right| B_1^\dagger B_2^\dagger B_2 B_1 \left|b_{in}\right\rangle \delta_{14}\delta_{23} \\
& = \alpha^2 g^{(2)}_{\mathrm{B,in}} \int \mathrm{d}t \int \mathrm{d}_{12}\, G^*_{sbt1} G^*_{sbt2} G_{sbt2} G_{sbt1} \\
& = 2\alpha^2 \int \mathrm{d}t \int \mathrm{d}_1\, \left|G_{sbt1} G_{sbt1}\right|^2 \\
& \equiv 2\alpha^2 \mathcal{G}_{sb},
\end{aligned}
$$

since the spin wave has thermal statistics and $g^{(2)}_{\mathrm{B,in}} = 2$. The cross terms containing S and A, $\langle S \times A \rangle$ consist of four terms:

$$
\begin{aligned}
\langle \mathrm{S} \times \mathrm{A} \rangle &= \int \mathrm{d}t \int \mathrm{d}_{1234} \left(G^*_{sst1} G^*_{sa^\dagger t2} G_{sa^\dagger t3} G_{sst4} \langle S_1^\dagger A_2 A_3^\dagger S_4 \rangle + G^*_{sst1} G^*_{sa^\dagger t2} G_{sst3} G_{sa^\dagger t4} \langle S_1^\dagger A_2 S_3 A_4^\dagger \rangle \right) \\
&+ \int \mathrm{d}t \int \mathrm{d}_{1234} \left(G^*_{sa^\dagger t1} G^*_{sst2} G_{sa^\dagger t3} G_{sst4} \langle A_1 S_2^\dagger A_3^\dagger S_4 \rangle + G^*_{sa^\dagger t1} G^*_{sst'2} G_{sst3} G_{sa^\dagger t4} \langle A_1 S_2^\dagger S_3 A_4^\dagger \rangle \right) \\
&= N_{\mathrm{in}} \int \mathrm{d}t \int \mathrm{d}_{1234} \left(G^*_{sst1} G^*_{sa^\dagger t2} G_{sa^\dagger t3} G_{sst4} \delta_{14}\delta_{23} + G^*_{sst1} G^*_{sa^\dagger t2} G_{sst3} G_{sa^\dagger t4} \delta_{13}\delta_{24} \right) \\
&+ N_{\mathrm{in}} \int \mathrm{d}t \int \mathrm{d}_{1234} \left(G^*_{sa^\dagger t1} G^*_{sst2} G_{sa^\dagger t3} G_{sst4} \delta_{13}\delta_{24} + G^*_{sa^\dagger t1} G^*_{sst2} G_{sst3} G_{sa^\dagger t4} \delta_{13}\delta_{24} \right) \\
&= 4N_{\mathrm{in}} \int \mathrm{d}t \int \mathrm{d}_{12}\, \left|G_{sst1}\right|^2 \left|G_{sa^\dagger t2}\right|^2 \\
&= 4N_{\mathrm{mem}} \times N^{AS}_{\mathrm{spont}},
\end{aligned}
$$

and similarly

$$
\begin{aligned}
\langle \mathrm{S} \times \mathrm{B} \rangle &= 4N_{\mathrm{mem}} \times N^{P}_{\mathrm{spont}} \\
\langle \mathrm{A} \times \mathrm{B} \rangle &= 4N^{AS}_{\mathrm{spont}} \times N^{P}_{\mathrm{spont}}.
\end{aligned}
$$

Putting this all together we have:

$$
\iint \mathrm{d}t\, \mathrm{d}t' \left\langle S^\dagger_{t'} S^\dagger_t S_t S_{t'} \right\rangle = N_{\mathrm{in}} g^{(2)}_{\mathrm{in}} \mathbb{G}_{ss} + 2\mathbb{G}_{sa^\dagger} + 2\alpha \mathbb{G}_{sb} + 4\left(N_{\mathrm{mem}} N^{AS}_{\mathrm{spont}} + N^{AS}_{\mathrm{spont}} N^{P}_{\mathrm{spont}} + N_{\mathrm{mem}} N^{P}_{\mathrm{spont}} \right)
$$

and recalling from Equation A.4 that $\int \mathrm{d}t \left\langle S_t^\dagger S_t \right\rangle = N_{\mathrm{mem}} + N^{AS}_{\mathrm{spont}} + N^{P}_{\mathrm{spont}}$, this gives:

$$
\begin{aligned}
\iint \mathrm{d}t\, \mathrm{d}t' \left\langle S^\dagger_{t'} S^\dagger_t S_t S_{t'} \right\rangle &= 2\left(\int \mathrm{d}t \left\langle S_t^\dagger S_t \right\rangle \right)^2 - N^2_{\mathrm{in}} (2\eta^2 - g^{(2)}_{\mathrm{in}} \mathcal{G}_{ss}) - 2((N^{AS}_{\mathrm{spont}})^2 - \mathcal{G}_{sa^\dagger}) \\
&\quad -2((N^{P}_{\mathrm{spont}})^2 - \alpha^2 \mathcal{G}_{sb})
\end{aligned}
$$

where $\eta = N_{\mathrm{mem}}/N_{\mathrm{in}}$. Therefore

$$g^{(2)}_{\text{out,sig}} = \frac{\iint \mathrm{d}t\,\mathrm{d}t'\,\langle S^\dagger_{t'} S^\dagger_t S_t S_{t'}\rangle}{\left(\int \mathrm{d}t\,\langle S^\dagger_t S_t\rangle\right)^2}$$
$$= 2 - \frac{N^2_{\text{in}}(2\eta^2 - g^{(2)}_{\text{in}}\mathcal{G}_{ss}) + 2((N^{AS}_{\text{spont}})^2 - \mathcal{G}_{sa^\dagger}) + 2((N^{P}_{\text{spont}})^2 - \alpha^2\mathcal{G}_{sb})}{(\eta N_{\text{in}} + N^{AS}_{\text{spont}} + N^{P}_{\text{spont}})^2}. \quad \text{(A.5)}$$

We now assume that, if there is no input signal, the noise due to spontaneous Raman scattering would give thermal output statistics i.e. $g^{(2)}_{\text{out}}(N_{\text{in}} = 0) = 2$. If $N^P_{\text{spont}} = 0$ and the thermal noise is only due to spontaneous anti-Stokes scattering:

$$g^{(2)}_{\text{out,sig}} = 2 - \frac{2((N^{AS}_{\text{spont}})^2 - \mathcal{G}_{sa^\dagger})}{(N^{AS}_{\text{spont}})^2} = 2$$

and so we can approximate $\mathcal{G}_{sa^\dagger} = (N^{AS}_{\text{spont}})^2$. Similarly if $N^{AS}_{\text{spont}} = 0$ and we only have thermal noise due to Raman scattering from the unpumped population:

$$g^{(2)}_{\text{out,sig}} = 2 - \frac{2((N^{P}_{\text{spont}})^2 - \alpha^2\mathcal{G}_{sb})}{(N^{P}_{\text{spont}})^2} = 2$$

and $\mathcal{G}_{sb} = (N^P_{\text{spont}})^2/\alpha^2$. With these approximations, Equation A.5 then simplifies to

$$g^{(2)}_{\text{out,sig}} = 2 - \frac{N^2_{\text{in}}(2\eta^2 - g^{(2)}_{\text{in}}\mathcal{G}_{ss}))}{(\eta N_{\text{in}} + N_{\text{spont}})^2}$$

where $N_{\text{spont}} = N^{AS}_{\text{spont}} + N^P_{\text{spont}}$.

We treat the noise contribution from fluorescence as an incoherent sum of fields with N_F photons and $g^{(2)}_F$. The incoherent sum of fields 1 and 2 is given by [183]:

$$g^{(2)}_{12} = \frac{N^2_1 g^{(2)}_1 + 2N_1N_2 + N^2_2 g^{(2)}_2}{(N_1 + N_2)^2},$$

and hence

$$g^{(2)}_{\text{tot}} = 1 + \frac{N^2_S - N^2_{\text{in}}(2\eta^2 - g^{(2)}_{\text{in}}\mathcal{G}_{ss})}{N^2_{\text{tot}}} + \frac{N^2_F(g^{(2)}_F - 1)}{N^2_{\text{tot}}},$$

where photons produced in the signal mode are $N_{\text{S}} = \eta N_{\text{in}} + N_{\text{SRS}}$, and the total number of photons produced is $N_{\text{tot}} = \eta N_{\text{in}} + N_{\text{SRS}} + N_F$. This leads to our final expression for $g^{(2)}_{\text{out}}$:

$$\boxed{g^{(2)}_{\text{out}} = 1 + \frac{aN^2_{\text{out}} + 2N_{\text{SRS}}N_{\text{out}} + b}{(N_{\text{out}} + N_{\text{SRS}} + N_{\text{F}})^2}} \quad \text{(A.6)}$$

where

$$\begin{aligned} a &= g_{\mathrm{in}}^{(2)}\mathcal{G}_{ss}/\eta^2 - 1 \\ b &= N_{\mathrm{SRS}}^2 + N_{\mathrm{F}}^2(g_{\mathrm{F}}^{(2)} - 1). \end{aligned} \tag{A.7}$$

In the case where there is no fluorescence noise, ($N_{\mathrm{F}} = 0$), and the input is a single photon with $g_{\mathrm{in}}^{(2)} = 0$:

$$\begin{aligned} \mu_1 &= \frac{N_{\mathrm{spont}}}{\eta} \\ g_{\mathrm{out}}^{(2)} &= 2\left[1 - \frac{(\eta N_{\mathrm{in}})^2}{(\eta N_{\mathrm{in}} + N_{\mathrm{spont}})^2}\right]. \end{aligned}$$

To observe a non-classical output, $g_{\mathrm{out}}^{(2)} < 1$, we therefore require an input photon number, or heralding efficiency η_h, greater than:

$$\eta_h > (1 + \sqrt{2})\mu_1. \tag{A.8}$$

We see that μ_1 can give us an estimate of whether the noise is sufficiently low to retrieve a non-classical state from the memory, and by measuring $g_{\mathrm{out}}^{(2)}$ as a function of the input photon number, N_{in}, we can determine the relative contribution of different noise processes and predict $g_{\mathrm{out}}^{(2)}$ for a stored single photon.