CDT in Controlled Quantum Dynamics
Department of Physics

Imperial College
London

Optimisation and Applications of a Raman Quantum Memory for Temporal Modes of Light

Joseph H. D. Munns

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

Submitted for the degree of Doctor of Philosophy in Physics

August 2018

DECLARATIONS

I declare that this thesis, and the work presented within, are my own and was completed whilst a candidate for a research degree at Imperial College London.

The experimental results herein have been obtained either by me, or in collaboration with colleagues whom I have credited where appropriate. I have indicated which derivations are, to my knowledge, my own original work; and, have appropriately referenced the sources of all other results.

Oxford, August 2018

Joseph Hugh Deakin Munns

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non–Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work

ABSTRACT

Within any practical proposal for utilising quantum mechanics for information processing whether for the efficient computation of problems intractable by classical methods, or for secure communication; at some stage one must have the means to transfer quantum information between remote nodes of a network. For this, light is the obvious choice.

To realise this vision one requires the means to overcome the "scaling problem" intrinsic to many photonic-based quantum technologies due to probabilistic operations. It has been identified that optical quantum memories which facilitate the storage and retrieval of quantum states of light are an enabling technology. Another requisite technology to ensure the scalability of a quantum network is the means to interface dissimilar material nodes, which in practice means the translation of quantum light in bandwidth, frequency and temporal shape.

The first part of this thesis presents experimental, theoretical and numerical investigations of noise reduction in the Raman memory protocol in thermal caesium vapour, by means of a cavity. To do this, I develop a theoretical description of the cavity memory interaction, along with a model of the atom-cavity system to enable meeting the required resonance conditions. This is followed by a proof-of-concept experimental demonstration, showing suppression of noise in the retrieved state. To conclude this part, I investigate the optimisation of this system and provide a numerical framework for its design, and propose a route towards realising the Raman memory as a practical quantum memory.

The second theme is an exploration of the practical application of the Raman memory as an interface for temporal modes of light. I perform a preliminary investigation, and develop characterisation tools, to experimentally verify the modal structure of the memory interaction. This work provides the basis for deploying the Raman memory as a temporal-mode selective device for GHz bandwidth quantum states of light.

ACKNOWLEDGEMENTS

Thanks to many people! Yer'all great.

Firstly here's to Ian, Myungshik, and in particular Josh (J.N.).

Thanks to everyone whom I have worked alongside over the last four years without whom this would not have been possible: Dylan (D.J.S.) and Tessa (T.F.M.C.), for all you both taught me during the cavity memory experiment; Patrick (P.M.L) for the sage advice and pints; Benni (B.B.) for introducing me to temporal modes; and Sarah (S.E.T.) it's been great working with you on the memory tomography; also to Kris (K.T.K.) and Cheng (C.Q.). You have all been fantastic. And thanks to the rest of the Memories team, past and present. Thanks are also due to everyone who read parts of this document at varying stages of completion (J.N., D.J.S., P.M.L., B.B. and S.E.T.) and the encouragement and helpful comments – I couldn't have put this together otherwise.

As well Memories, the wider Ultrafast group is a great community, I've learned a lot and I'm glad to have been part of it. Thanks to Amir for bringing me into the sims fold, which has been an educational diversion, and to Henry for dreaming up the "solattice" project.

Thanks to the Imperial cohort, in particular the Oxford division. Cheers comrades!

Finally, here's to those outside of work who have helped me to get here. So, thanks to MCW, PVM, EMD...*et al.*!

AUTHOR PUBLICATIONS

Articles

Related Publications

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

"Theory of noise suppression 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, Dileep V. Reddy, M.G. Raymer & I.A. Walmsley
Phys. Rev. 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 J. Phys. **19** 063034 (2017)

"A Cavity–Enhanced Room–Temperature Broadband Raman Memory"
D.J. Saunders, J.H.D. Munns, T.F.M. Champion, C. Qiu, K.T. Kaczmarek, E. Poem, P.M. Ledingham, I.A. Walmsley, & J. Nunn
Phys. Rev. Lett. **116**, 090501 (2016)

"*In Situ* Characterisation of an Optically Thick Atom–Filled Cavity"
J.H.D. Munns, C. Qiu, P.M. Ledingham, I.A. Walmsley, J. Nunn, & D.J. Saunders
Phys. Rev. A **93**, 013858 (2016)

Other

"Soliton Synchronisation and Hopping in Kerr Soliton–Cavity–Networks"
J.H.D. Munns, I.A. Walmsley, & Y.H. Wen
In preparation

"A Room–Temperature Noise–Free Quantum Memory for Broadband Light"
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
Phys. Rev. A **97**, 042316 (2018)

"Broadband noise–free optical quantum memory with neutral nitrogen–vacancy centers in diamond"
E. Poem, C. Weinzetl, J. Klatzow, K.T. Kaczmarek, J.H.D. Munns, T.F.M. Champion, D.J. Saunders, J. Nunn, & I.A. Walmsley
Phys. Rev. B **91**, 205108 (2015)

Conference Papers

Oral Presentations

"Novel Interactions of Dissipative Kerr Solitons in Coupled Cavity Networks"
J.H.D. Munns, I.A. Walmsley & Y.H. Wen
EF-1.1 *CLEO/Europe–EQEC*, Munich, Germany, 2017

"Bad Cavities for Good Memories"
J.H.D. Munns, T.F.M. Champion, C. Qiu, P.M. Ledingham, D.J. Saunders, I.A. Walmsley, & J. Nunn
FW4D.7 *OSA Frontiers In Optics*, San Jose, USA, 2015

"Towards Noise Free Quantum Memories"
J.H.D. Munns, T.F.M. Champion, K.T. Kaczmarek P.M. Ledingham, D.J. Saunders, I.A. Walmsley & J. Nunn
9378-29 *SPIE Photonics West*, San Francisco, USA, 2015

Poster Presentations

"A High-Speed Noise-Free Optical Quantum Memory"
K.T. Kaczmarek, J.H.D. Munns, P.M. Ledingham, B. Brecht, S.E. Thomas, G.S. Thekkadath, O. Lazo-Arjona, E. Poem, A. Feizpour, D.J. Saunders, J. Nunn, & I.A. Walmsley
Quantum Computing Measurement & Control, Louisiana, USA, 2018

"Temporal-mode selection with a Raman quantum memory"
J.H.D. Munns, S.E. Thomas, K.T. Kaczmarek, P.M. Ledingham, D.J. Saunders, J. Nunn, B. Brecht & I.A. Walmsley
(selected for Rapid Fire Oral Presentation)
JW4A.16 *OSA Frontiers In Optics*, Washington D.C., USA, 2017

"Temporal Amplitude & Phase: Algorithmic Reconstruction via Time-domain Interferometry"
J.H.D. Munns, S. Park, B. Brecht & I. A. Walmsley
(selected for Rapid Fire Oral Presentation)
JW3A.34 *OSA Frontiers In Optics*, Washington D.C., USA, 2017

"Modelling Atom–Filled Optical Cavities for Enhanced Light–Matter Interaction"
J.H.D. Munns, S.E. Thomas, K.T. Kaczmarek, B. Brecht, P.M. Ledingham, D.J. Saunders, J. Nunn & I.A. Walmsley
Quantum Computing Measurement & Control, Singapore, 2016

"Suppression of Four–Wave Mixing in a Raman Memory"
J.H.D. Munns, T.F.M. Champion, C. Qiu, K. T. Kaczmarek P.M. Ledingham, D.J. Saunders, J. Nunn & I.A. Walmsley
QUTE–Europe Summer School, Gothenburg, Sweden, 2015

"Storage of Single Photons in a Raman Memory"
T.F.M. Champion, P.S. Michelberger, J.H.D. Munns, M.R. Sprague, K.T. Kaczmarek, M. Barbieri, X.M. Jin, D.G. England, W.S. Kolthammer, D.J. Saunders, J. Nunn & I.A. Walmsley
Quantum Information, Computing & Control – CQD CDT Summer School, Exeter, UK 2014

CONTENTS

Declarations . . . 2
Abstract . . . 3
Acknowledgements . . . 4
Author Publications . . . 5
List of Figures . . . 10
List of Tables . . . 12
Abbreviations, Symbols and Constants . . . 13

1 Introduction **15**
1.1 Architectures for Quantum Information Processing . . . 15
1.2 Photonics . . . 17
1.3 A Brief Overview of Quantum Memories . . . 18
1.4 Schemes & Platforms . . . 23
1.5 Thesis Outline . . . 27

2 The Raman Quantum Memory **28**
2.1 Introduction . . . 28
2.1.1 Noise: Spontaneous Four-Wave Mixing . . . 29
2.2 Equations of Motion . . . 30
2.3 Ideal Operation . . . 38
2.3.1 Memory . . . 39
2.4 Raman Selection Rules . . . 41
2.5 Summary . . . 42

I Suppression of Four–Wave Mixing in Raman Memories **43**

3 A Cavity Enhanced Raman Memory **44**
3.1 The Idea . . . 44
3.2 Theory of Noise Suppression . . . 45
3.2.1 Bad Cavity Limit . . . 48
3.2.2 Performance . . . 52
3.2.3 Theory Summary . . . 61
3.3 Cavity Engineering . . . 64
3.3.1 Cavity Memory Experiment . . . 64
3.3.2 Atoms In A Cavity . . . 67
3.3.3 Dispersion Model Applied to the Cavity Memory . . . 71
3.3.4 Numerics . . . 73
3.4 Implementation . . . 75
3.4.1 Experimental Details . . . 76
3.4.2 Memory Results . . . 79
3.4.3 Storage of Coherent States . . . 79

3.4.4 Noise Analysis . . . 83
3.5 Summary . . . 84

4 Cavity Memory Optimisation **85**
4.1 Introduction . . . 85
4.1.1 Memory Efficiency . . . 85
4.1.2 Memory Lifetime . . . 86
4.1.3 Noise Reduction . . . 87
4.2 Buffer Gas (1): Atomic Diffusion Measurements . . . 87
4.2.1 Method . . . 88
4.2.2 Experimental Details . . . 89
4.2.3 Implications . . . 92
4.3 Buffer Gas (2): Quenching . . . 96
4.4 Numerical Parameter Search . . . 97
4.4.1 Cavity Raman Memory . . . 97
4.4.2 Numerical Implementation . . . 99
4.5 Summary . . . 107

5 Conclusions Part I **108**
5.1 Summary . . . 108
5.2 Wider Landscape . . . 109
5.3 Moving Forward . . . 110
5.3.1 Improved Cavity Design . . . 110
5.3.2 Alternative Approaches to Noise Suppression: "Magic Λ" . . . 110
5.3.3 Temporal Modes . . . 111

II Temporal Modes of Light **112**

6 Temporal Modes of Light: An Introduction **113**
6.1 Making use of Temporal Shapes in Light . . . 113
6.1.1 Temporal Modes for QIP . . . 113
6.1.2 Implementing TM encoding . . . 114
6.1.3 Enhanced Light–Matter Interaction . . . 115
6.2 Devices for Temporal Modes . . . 115
6.2.1 Quantum Pulse Gates . . . 117
6.2.2 Single Mode Quantum Memories . . . 117
6.3 Outlook . . . 119

7 Mode–Selectivity in a Raman Memory **120**
7.1 Introduction . . . 120
7.2 Concepts . . . 121
7.2.1 Tomographic Analysis . . . 121
7.2.2 Modelling the Memory Interaction . . . 122
7.3 Experiment . . . 124
7.3.1 State Preparation . . . 124
7.3.2 Memory infrastructure: Modifications . . . 128
7.3.3 Preliminary Conclusions . . . 131
7.4 Reconstruction of the Green's Function . . . 133
7.4.1 Informationally Incomplete Tomography . . . 133
7.4.2 Mapping to Memory Experiment (1): Theory . . . 134
7.4.3 Mapping to Memory Experiment (2): Practice . . . 135
7.4.4 Performance Benchmarking . . . 136
7.5 Results . . . 137

7.5.1 Experimental Data . . . 137
7.5.2 Improving Measurement Choices . . . 139
7.5.3 Conclusions . . . 142
7.6 Summary . . . 147

8 Conclusions Part II **149**
8.1 Summary . . . 149
8.1.1 Mode Selectivity . . . 149
8.1.2 Mode Reconstructions . . . 150
8.2 Future . . . 150
8.2.1 "Stark Compensated Cavity–Magic–Raman" Memory . . . 151
8.2.2 Photon Source Purification . . . 152
8.2.3 Mode Conversion . . . 152
8.3 Outlook . . . 153
8.4 Final Thoughts . . . 153

Back Matter **154**

Bibliography **155**

A Experimental Control Infrastructure **169**
A.1 Cavity Memory Experiment: Control Systems . . . 169
A.2 Memory Optimisation Experiments . . . 174
A.3 Pulse Carving . . . 174

B TEA+PARTI **177**
B.1 Pulse Shaping Infrastructure . . . 177
B.2 TEA+PARTI: Basic Idea . . . 178
B.3 TEA+PARTI in Practice . . . 184
B.4 Applications . . . 186
B.5 Conclusions . . . 189

C Optimisation of AWG Electrode Voltages **190**
C.1 Downhill Simplex Search . . . 190
C.2 Future . . . 193

D Reconstruction Algorithm **196**
D.1 Informationally Incomplete Maximum Likelihood Tomography . . . 196

E Mode Selectivity Measurements & Numerics **200**
E.1 Memory Measurements . . . 200
E.2 Quadrature Methods . . . 203
E.3 Reconstruction Basis . . . 204
E.4 Tomography Algorithm . . . 204

LIST OF FIGURES

1.1 Bell-state measurement circuit . . . 20
1.2 Temporal multiplexing with quantum memories . . . 21

2.1 Atomic energy level schematic . . . 30

3.1 Principle of noise suppression in a cavity . . . 45
3.2 Ring cavity geometry . . . 47
3.3 Schematic of ring cavity implemented experimentally, and optical pumping scheme . . . 65
3.4 Doppler-free spectrum of Cs D_2 line . . . 67
3.5 Simple "atom filled cavity" schematic . . . 68
3.6 Number density of caesium atoms as a function of temperature . . . 72
3.7 Pumping efficiency and atom number density as a function of cell temperature . . . 74
3.8 Typical transmission spectrum, near atomic resonance with different number densities 76
3.9 Measured & fitted cavity response locked to signal frequency at varied cell temperatures 77
3.10 Cavity memory experimental setup . . . 78
3.11 Typical plot of time tagger data for memory experiments . . . 80
3.12 Cavity–Raman memory performance . . . 82
3.13 Noise metrics for the cavity–Raman memory . . . 83

4.1 Cartoon of pump/probe measurement of diffusion . . . 88
4.2 Sequence of images from which atomic diffusion rate can be extracted . . . 90
4.3 1D cross section of optical depth profile . . . 91
4.4 Extraction of D_0 for three buffer gas species. . . . 93
4.5 Projected lifetime achievable with N_2 buffer gas, varying pressure and beam waist . . . 95
4.6 Simulated cavity spectra meeting triple resonance varying detuning and Cs density keeping intracavity loss constant . . . 100
4.7 Parameter search investigating asymtotic efficiency achievable in $M = 0$ cavity . . . 101
4.8 Parameter search of η, SNR, $g^{(2)}$ in $M = 0$ cavity, with and without optimal mode matching with constant intracavity loss . . . 102
4.9 Asymptotic efficiency as incoupler/intracavity loss is varied . . . 105
4.10 Efficiency and $g^{(2)}$ allowing free variation of detuning and caesium density, with fixed Ω 106

5.1 Alternative memory configuration . . . 111

6.1 Optimal control field for HG signals . . . 116
6.2 Schematic of QPG operation . . . 118
6.3 Cartoon of TM selective operations . . . 119

7.1 Measured carver electrode signals, and inferred optical fields . . . 125
7.2 Optimistic estimate of MUB state preparation fidelity . . . 127
7.3 Optical intensity of the MUB probe states . . . 128
7.4 Memory reconstruction experiment . . . 129

7.5 Normalised efficiencies for storage of MUB states . . . 130
7.6 Simulation of memory performance . . . 131
7.7 Measured memory efficiency for timescanned probe states in 2D . . . 132
7.9 Comparison of reconstruction on experimental data with model . . . 138
7.10 Reconstruction on simulated efficiencies with ideal probe states in 3D . . . 140
7.11 Reconstruction on simulated efficiencies with width mismatched probe states . . . 141
7.12 Interaction matrix eigenmode evolution as a function of Rabi frequency amplitude . . . 143
7.13 Mode weighting as a function of Rabi frequency amplitude . . . 145
7.14 Comparison of extracted effective mode number directly/reconstructed from simulated matrix . . . 146

A.1 First implementation of cell temperature control and monitoring . . . 170
A.2 Main loop of Arduino locking code . . . 173
A.3 Measurement of cavity mechanical stability . . . 174
A.4 Bulk optical depth calibration and measurement . . . 175

B.1 Fibre Mach-Zehnder for arbitrary pulse carving . . . 178
B.2 Schematic for the interference measurement in TEA+PARTI. . . . 179
B.3 fourier filtering of signal . . . 180
B.4 In situ shear frequency calibration measurement . . . 183
B.5 Experimental setup of the TEA+PARTI interferometer. . . . 185
B.6 TEA+PARTI reconstruction on experimental data . . . 187

C.1 Parametrisation of search vector . . . 192
C.2 Nelder-Mead iteration process . . . 193
C.3 Evolution of the downhill simplex algorithm . . . 194
C.4 Amplifier outputs before/after downhill simplex search . . . 195

E.1 Memory data for carved input signals . . . 201
E.2 Detail of time scanned 2D memory measurements . . . 202
E.3 Comparison of different quadrature methods . . . 203
E.4 Coordinate transformation from “lab time” to HG basis . . . 205
E.5 Typical trajectory of the tomography algorithm convergence . . . 205

LIST OF TABLES

1.1 State of the memories community 27

2.1 Labels for the respective detunings between the atomic and optical frequences. 33

3.1 Autocorrelation statistics of the Raman memory 56
3.2 Asymptotic limits 62
3.3 Comparison of parameters extracted via automated and manual fitting routines 75

4.1 Pressure broadening coefficients for various buffer gas candidates 88
4.2 Extracted diffusion parameters for different buffer gases 92
4.3 Target beam waist for different ring cavity optics 94
4.4 Estimated memory lifetime for different cavity memory configurations 96
4.5 Suppression factor, x, for potential cavity configurations 104
4.6 Relation between key experimental parameters and resulting figures of merit 107

7.1 MUB states for a qubit 121

ABBREVIATIONS, SYMBOLS AND CONSTANTS

List of Abbreviations

AOM Acousto–Optic Modulator

APD Avalanche Photo–Diode

BS Beam–Splitter

CW Continuous Wave

DWDM Dense Wavelength Division Multiplexing

ECDL External Cavity Diode Laser

EOM Electro–Optic Modulator

FMZI Fibre Mach–Zehnder Interferometer

FWM Four–wave mixing

HG Hermite–Gauss (polynomials)

LG Legendre–Gauss (quadrature)

MMF Multi-Mode Fibre

PBS Polarising Beam–Splitter

POVM Positive Operator Value Measure

QM Quantum Memory

QPG Quantum Pulse-Gate

SMF Single–Mode Fibre

SPDC (Spontaneous) Parametric Down–Conversion

SPIDER Spectral Phase Interferometry for Direct Electric–field Reconstruction

TEA+PARTI Temporal Electric–field Amplitude & Phase: Algorithmic Reconstruction via Time–domain Interferometry

Ti:Sa Titanium Sapphire oscillator

List of Symbols and Constants

Symbolic

$*$ ..... Complex conjugate

$\dagger$ ..... Hermitian conjugate

∇ ..... Nabla gradient operator

$\otimes$ ..... Tensor product

Numeric

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

Greek

Δ ..... Frequency detuning

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

$\Theta(\cdot)$ ..... Heaviside step function

$\delta(\cdot)$ ..... Dirac delta

$\delta_{i,j}$ ..... Kronecker delta

ϵ ..... Normalised time

ε ..... Dielectric permittivity

ε_0 ..... Permittivity of free space

μ_0 ..... Magnetic permeability of free space

μ_1 ..... Metric for noise level

ω ..... Angular frequency

Latin

$\mathcal{F}[\cdot]$ ..... Fourier transform of $\cdot$

$\mathcal{H}$ ..... Hilbert space

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

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

$T_{i\in\{1,2\}}$ ..... Decay constant ($i = 1$: longitudinal, $i = 2$: tranverse)

$\mathrm{Tr}[\cdot]$ ..... Trace of $\cdot$

c ..... Speed of light

e ..... Euler's number

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

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

i ..... Imaginary unit

k_B ..... Boltzmann constant

CHAPTER 1

INTRODUCTION

As well as quantum mechanics being a fundamental theory with deep, and sometimes unsettling, consequences on the nature of reality, the predictions from its laws make it one of the most successful physical theories. These consequences extend even to our understanding of logic, and in the last few decades there has been a drive to understand if and how we may use this to do useful things.

The technological applications[1] of quantum information science covers a diverse range: from simulations (simulating quantum systems with quantum systems) [1], quantum metrology and sensing [2] and secure communication [3], up to a fully fledged quantum computer [4]. A relatively recent addition to this field applies this approach to quantum machine learning[2] such as pattern finding [5]. This list goes on, and the field will no doubt continue to grow!

In all of these cases, the elements necessary for computation and simulation rely upon ones ability to perform sequential operations on the elementary components of the system.

1.1 Architectures for Quantum Information Processing

The most direct analogue to classical computation is the circuit model for quantum computation. This describes how an algorithm may be built up from a series of elementary logic gates acting on one, or between two quantum bits (qubits), enabling "universal quantum computation" [6]. The development of quantum algorithms which promise enhancements over their classical counterparts is a growing field [7–9] – where exploiting quantum mechanics, in particular interference and entanglement[3] – enables dramatic speed up[4] compared to the best known classical algorimthms, for certain problems.[5]

Beyond the circuit model, there are other promising architectures such as cluster states and measurement-based (or "one-way") quantum computation first proposed in [11], in which a highly entangled cluster state is consumed as a resource. But, once it is generated, universal computation[6] can

[1] Without considering the more foundational aspect of quantum information.

[2] This is based upon linear matrix operations in high dimensional vector spaces – a problem to which quantum mechanics is very well suited!

[3] The structure of many of these algorithms is based upon engineering interference such that the "correct answer" is enhanced, while others destructively interfere, followed by quantum phase estimation [10].

[4] The *efficiency* of an algorithm is quantified in terms of the scaling with with respect to the problem size. A change in scaling from exponential time to polynomial takes a problem from being intractable to being a solveable task.

[5] Indeed some quantum algorithms are hybrid and call a classical subroutine [e.g. 9], for parts of the computation where no quantum enhancement is found.

[6] Provided that the structure of the graph is sufficiently interconnected (for instance, due to experimental imperfection, see e.g. [12]) such that the desired algorithm can be mapped to a measurement sequence on the graph [13] – universal computation is not possible with a 1D chain, but a fully connected 2D cluster would be universal.

be implemented via single qubit operations and measurements[7] with classical feed-forward. Schemes for scalable quantum computation using cluster states include proposals such as the generation of 3D graph states for photonic implementation with a continuous stream of nodes added to the cluster state [15], where once the resource state has been generated, the implementation of single qubit operations on each of the photonic nodes are[8] straightforward to implement.

Similarly an area of interest for quantum processors[9] is the simulation of materials – for instance the problem of finding the ground state of complex molecules; which quickly becomes intractable with classical computation as the size of the system increases, and approximations or coarse graining approaches employed limit the achievable accuracy. There are two principal categories for quantum simulation [16]. Firstly, adiabatic simulation whereby the unknown ground state of some Hamiltonian is found by preparing the system in some easily prepared ground state of an initial Hamiltonian, and adiabatically tuning the Hamiltonian to the the final state [17, 18] where this ground state, could be mapped to the solution of some other computationally interesting problem [19]. The alternative approach, more related to the gate model of quantum computation, is based upon approximating the Hamiltonian by a sequence of small operations, or "Trotterisation", followed by quantum phase estimation [20, 21].

The applications of quantum information processing (QIP) are many and diverse, and different schemes will naturally lend themselves to different physical architectures; but at the basic level, these all rely on our ability to control – that is: prepare; manipulate; and read out – isolated quantum systems.

Over the recent years there has been much drive in moving forward our ability to experimentally perform these elementary operations in different potential media, and to concatenate these elementary operations to build up to larger scale computations in a range of different media, such as in superconducting circuits, trapped atoms or ions, nitrogen-vacancy centres or quantum dots [22–26].

To compare the performance of different platforms, one must consider a range of figures of merit. Perhaps most intuitively, the fidelity of gate operations, both upon single qubit but also for two qubit entangling operations. But also one should consider the relevant timescales of interacting with a given system – i.e. the coherence time of the qubit, and the time taken to implement various operations.

Both of these impact on the potential for *scalability* of some platform for quantum computation; primarily in terms of its usefulness as a computer where one should aim to surpass a threshold for a fault tolerant scheme [27, 28], whereby errors due to decoherence or experimentally imperfect gates can be corrected with certainty by syndrome measurements;[10] but also the scaling of these metrics when one considers the more experimental overheads for concatenating many operations.

There exist proposals to build up large scale quantum computers in a single system, such as trapped ion computing using microfabricated chips on which many ions may be manipulated and shuttled [29, 30] – or in a constrained geometry making use of collective motional excitations for ions, and performing "hide" operations to selectively address qubit(s) in the chain as was used for the demonstration of deterministic teleportation (within the same ion trap) [31]. To further increase the

[7] In their original formulation [11], the authors show that Clifford group gates, plus *arbitrary* $z-$rotations ($\exp[-i\hat{\sigma}_z\theta/2]$), and Pauli measurement can be used to implement *any* single qubit rotations on the cluster. This is necessary to surpass the Gottesman-Knill theorem [14] which states that a process is efficiently simulable on a classical computer if it consists only of Clifford group operations (the CNOT, Hadamard, and $\pi/2-$Phase gates), and measurement of the Pauli operators.

[8] ...in principle, once one has the means synchronising these operations and being able to implement classical feed-forward based upon the outcome of probabilistic operations.

[9] Not necessarily a fully fledged quantum computer.

[10] Note that these schemes generally necessitate encoding a single logical qubit in multiple physical qubits.

size of the computation – and adding more qubits – then requires building larger and larger traps, and increases the experimental overheads.

An alternative approach is to consider a hybrid architecture in which one may exploit the strengths of different platforms and integrate these together – for instance performing deterministic gates in some matter system containing a few qubits, and linking different nodes[11] with photons – for such hybrid systems one requires a robust interface between light and matter.

This lends itself to a modular – therefore readily scalable – approach, for instance in trapped ions (e.g. [32]), making use of the fast (e.g. [33]) and high fidelity (e.g. [34]) gates, which combined with realistic lossy photonic links, should be sufficient to enable freely-scalable fault tolerant quantum computation [35] – in which scheme, the authors consider an end-to-end photon detection probability of 2%.

This discussion naturally leads to the idea of quantum networks [36]. With a modular approach, the network is naturally expandable: as more nodes may be added, resilient: should one node fail, and can simultaneously exploit the strengths of various different systems and architectures. The networked concept also brings with it new applications, for example to enable the distribution of computational tasks, via a cloud quantum computing scheme [37], where a user with limited resources, may remotely commit tasks to a remote quantum computer which is "blind" to the details of the query.

One of the primary requirements to achieve this is the ability to form photonic links between each of these nodes. Quantum memories are a way to achieve this [38].

Furthermore, in order to simultaneously exploit the strengths of various platforms – or to interface non-identical devices in the solid state, each subject to a different local environment – this necessitates being able to interface between these. In general these systems will operate at dissimilar wavelengths and bandwidths. This means that another requisite for building up these networks is the means to coherently "translate" the signal in order to communicate between different devices.

As optical links are necessary to join the nodes of a quantum network, we turn our attention now to discuss photonics within QIP.

1.2 Photonics

Historically, photonics has often been the platform for fundamental experiments in foundations of quantum mechanics. For instance, to choose a recent example, the series of experiments in different groups seeking to close various loopholes in violating the Bell inequality [39–41].

One reason for the appeal of photonics is that photons themselves are weakly interacting with their environment – and, there is no thermal noise process at optical frequencies.[12] Equally the technical simplicity with which linear operations may be applied, polarisation rotated or pulses delayed.

Furthermore, as with classical communication, photons are ideal for transporting information, by definition *at the speed of light*. However single photons are limited to attenuation lengthscale on the order of a few tens of kilometres in fibre, before the photon is lost [42]. Additional strategies are needed to overcome this limit, and be able to establish a large scale "quantum internet".

In terms of building up large scale quantum states of light, the demonstration of 10 photon entanglement in 2016 [43] (via fusing five successively generated photon pairs, with success heralded by 10-fold coincidences at the output) is a promising result. However, this also highlights some of the

[11] Potentially in different systems.
[12] Intuitively seen if one compares the energy scales $k_B T \sim \hbar\omega$.

challenges of photonics as a platform. This record for photons is compared with approaches in other systems, such the demonstration of a coherent 51 qubit quantum simulator [44], or industrial efforts – typically based upon networks of superconducting qubits – such as the recently announced 72 qubit quantum processor [45].

1.2.1 Challenges and Components

Optics provides the basis for distributing information over large distances. In contrast to classical communication with light, loss is a critical factor for long-range communication with quantum states of light. Classical signals may be arbitrarily amplified, and may be robustly distributed over large distances via successive amplification, without any detriment to the signal. For the transmission of quantum information one is subject to the *no cloning theorem* [46, 47], and without any remedy a single photon may only cover a limited distance before it is lost. Whilst QKD protocols may not suffer in terms of security in the presence of loss, this will necessarily limit the rate at which information can be shared. In the next section we discuss the application of quantum memories as repeaters for the transmission of quantum states.

One of the features which makes photons ideal for the transmission of quantum information, is also one of their major drawbacks – it is hard to get photons to interact with one another, and conditional operations on pairs of bits, i.e. logic gates, are necessary in order to build up universal computation.

The Knill, Laflamme & Milburn (KLM) scheme [48] provides a framework for universal quantum computation which is built up from probabilistic gates based upon linear optical elements and measurement.[13] Despite the inherently probabilistic nature, the basic ingredients are technically simple to implement and the field has continued to attract interest in developing practical schemes, improving performance and correcting for errors [49]. One of the main limitations to scaling such architectures is understood to be the switching loss [50].

Arguably more desireable, but at the same time, more challenging to implement are schemes to implement deterministic single-photon nonlinearities for instance via quantum dots [51, 52]; Rydberg interaction [53], or engineering more traditonal non-linear crystal media [54]. Implementing this with single photons still remains a challenge.

1.3 A Brief Overview of Quantum Memories

A quantum memory (QM) can be broadly defined as a system which is "capable of storing a quantum state"[14] [42]. A subset of these are optical quantum memories, which specifically concerns the storage of quantum states of light; although in the following we will use the terms interchangeably.

This first definition captures the essence that the memory interaction is a coherent process, in that the state is reversibly mapped from one medium to another, *without disrupting the quantum state* for instance by measurement. An additional ingredient to consider is the manner of retrieval – for some purposes a "memory" which implements a pre-programmable delay might be sufficient; however, a more general-purpose memory enables read out "on demand" – that is that the time of read out may be chosen *after* the state is stored.

[13] A form of linear optical quantum computation (LOQC).

[14] Where in order to preserve a quantum state, a quantum memory must enact a linear map between the signal and storage modes.

Quantum memories have many applications, on a variety of scales and within a range of different architectures for QIP. For instance as local nodes for distributing quantum information over a large scale distributed quantum network; or for local synchronisation for processes at some quantum computing node.

In either case, we generally seek to employ a memory to enable the synchronisation of probabilistic events, such that states corresponding to successful outcomes can be held – whilst waiting for other processes to complete – and retrieved at some later time, *unspecified at the time of storage*.

1.3.1 Local & Distributed Synchronisation

1.3.1(a) Distributed: Quantum Repeaters

In building up a large scale quantum network, one requirement is to be able to generate *shared entanglement* between remote parties.[15] Shared entanglement can then be consumed as a resource which gives quantum communication protocols their power over classical counterparts.[16]

Scaling up quantum communication protocols to a global scale has received much attention recently following the recent demonstrations implementing two different QKD protocols from satellite-to-ground [58, 59]; which fits into a proposal for building up a global scale network of satellite mounted photon pair sources, and quantum repeaters at the ground stations [60].

In practice, building up a large scale quantum network will involve integrating quantum communication into existing optical fibre networks. When transmitting a photonic qubit between distant nodes, for instance down optical fibre, as the distance increases, the probability of loss and, for instance, depolarisation scales exponentially. Thus both the number of trials necessary to successfully share a photon increases (due to said loss) *and* the fidelity of the received state decreases (due to other channels). Even if purification protocols are employed, there is a minimum threshold of initial fidelity needed.

These together quickly place a practical upper limit on the distances over which a photonic qubit may be shared.

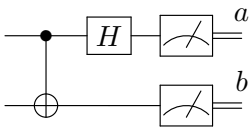
The quantum repeater, proposed by Briegel *et al.* in 1998 [61], provides the means of distributing entanglement over extended distance by subdividing a long quantum channel into segments over which an entangled pair may be safely shared, performing sequential entanglement swapping operations [62] to "fuse" adjacent segments of the channel together.

The process of entanglement swapping requires a Bell-state measurement (BSM) to be performed upon photons arriving from neighboring links. With only linear optics and photodetection, this can only be a *probabilistic* process [63]. Figure 1.1 shows a circuit to implement a Bell-state measurement: this requires a CNOT operation between the two modes, which cannot be achieved deterministically with linear operations, one requires single photon interactions which, as discussed at the start of this section, remains a challenge.

Memories are thus a vital component in the repeater scheme to enable synchronisation between different nodes and to allow feed forward via classical communication of the BSM results. Furthermore, since the measurement induced fusion between nodes is probabilistic it is advantageous for the

[15] Whilst discussing long-range communication, one should also consider the possibility of storing a quantum state, and physically transporting the "quantum memory stick", which becomes possible if the memory has sufficient lifetime. This is not as unrealistic as it might initially seem [55, 56]; although that memory stick would have to remain at cryogenic temperatures.

[16] Although, entanglement is not, of itself, necessary for security – as in the BB84 protocol [57].



input state	→	a	b
$\|00\rangle + \|11\rangle$	→	0	0
$\|01\rangle + \|10\rangle$	→	0	1
$\|00\rangle - \|11\rangle$	→	1	0
$\|01\rangle - \|10\rangle$	→	1	1

Figure 1.1: (left) Circuit to implement a Bell-state measurement; (right) corresponding truth table.

memories to have multi-mode capacity [64], such that multiple channels can be attempted in parallel, and the successful modes then released onto the stage on the network. In this way memories enhance the success rate of the distribution of entanglement.

1.3.1(b) Local: Synchronisation of Probabilistic Operations

This heading encompases a range of operations, such as the synchronisation of probabilistic logic gates implemented in linear optical computation (LOQC) schemes [48] constituting a system in which states may be temporarily stored in a "buffer memory". To begin at the start of the process, this encompases the generation of single photons via probabilistic means – such as spontaneous parametric downconversion (SPDC) in a non-linear cystal – which we take as an example to illustrate how QMs can be employed for local synchronisation operations.

Ideally, one would have access to a deterministic photon source – that is, a device which will produce an identical single photon (just one), on demand, every time. Indeed, there has been much progress in the development of deterministic photon sources for instance single atoms [65], single ions [66] – or artificial systems such as quantum dots[17] [69].

Historically the standard technique to produce single photons for photonic QIP is to use SPDC. This is a process where one high energy pump photon is downconverted (preserving energy and momentum) into a pair of correlated daughter photons with lower energy. Since the photons are produced in pairs, detection of one photon heralds the presence of the other photon. However, the probability p of downconversion is necessarily low, limited by the tolerable rate of multi-pair contributions, and therefore the rate of single pair production per pulse is low,[18] $p \lesssim 0.1$ [71].

In order to increase the probability of heralding a photon in the desired spatial mode, various multiplexing strategies could be adopted. The general scheme is that multiple SPDC sources are pumped in parallel – in some degree of freedom, and only the succesful heralded signals routed towards the output. One obvious solution is to use spatial multiplexing (e.g. refs [71, 72]), which makes use of many probabilistic sources operated in parallel and the outputs actively switched to select N sources in which a herald photon has been detected. However, this requires a large number of identical sources in order to be efficient, and it also requires $\log_2(N)$ switches to route the photon into the desired output [71]. This is a challenging task from an engineering standpoint. A more recent approach, less sensitive to switching loss,[19] is frequency multiplexing [73], in which photon pairs are generated over a range of frequencies and, based upon information about the frequency of the heralding photon, the frequency of the idler may be appropriately converted into a single output mode.

[17] Where these are also appealing for the promise of deterministic production of photonic cluster states [67], although a numerical study of the emission from localised light sources in micro-optical structures recently report a maximal collection efficiency limited to $\sim 30\%$ [68].

[18] And, the maximum probability of single pair production is 25% – in the ideal case, with zero losses and photon number resolving detectors [70].

[19] Since the switching depth is independent of the number of modes multiplexed.

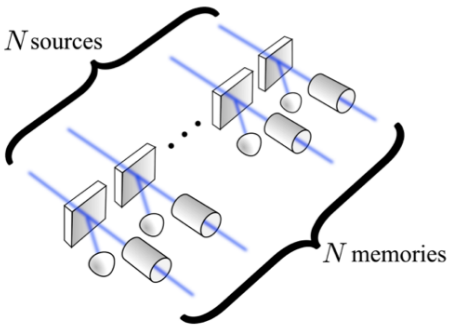


Figure 1.2: Schematic of a temporal multiplexing set up synchronised to increase the yield of N single photons from N SPDC photon sources. After ref. [74].

To reduce this scaling problem, another approach is to use a memory that can implement a temporal multiplexing strategy, in which fewer SPDC sources are pumped repeatedly and the single photons are stored as they are heralded [74]. The photons are then simultaneously released once the required number of photons have been prepared.

Of course, hybrid strategies incorporating both temporal and spatial multiplexing could help reduce the resource overheads of solely employing a single approach [75].

To highlight how a system of quantum memories can be used to enhance multiphoton rates, an example of a simple scheme is depicted in fig. 1.2, in which we consider a protocol requiring N coincident single photons from N SPDC sources. Without a memory present, all N sources must fire simultaneously, whereas with a memory single photons may be stored as soon as they are generated, and later read out once the N^{th} photon is heralded. Such a scheme illustrates how introducing a memory can reduce the overheads required to produce an N-fold coincidence of single photons, for increasing N. For example, as shown in ref. [74], for conceivable parameters, this could decrease the waiting time for a 10 photon state from ∼3 years to ∼0.01 seconds. Quantum memories would therefore prove a valuable enabling technology for the future of photonic QIP.

1.3.2 Figures of Merit

There are many and varied metrics for quantifying the performance of a memory, which depend upon the specific application, such as the two considered above. However the general constraints to be considered for a potential memory system used for these tasks include efficiency, storage time (along with the ability to read-out on demand), and importantly the preservation of the state stored [42].

For local synchronisation of local, *probabilistic*, applications, a key figure of merit is the product ηB [74] where η is the memory efficiency and B is the time-bandwidth product (TBP), defined as $B = \tau\delta$, with τ the memory lifetime and δ the memory acceptance bandwidth. The TBP may be interpreted as the ratio of the memory storage time to the repetition time of the memory i.e. the number of experimental clock-cycles within the memory lifetime. It has been shown that in these cases even a modest efficiency can greatly counter the "scaling problem" encountered when concatenating many probabilistic processes.

For long range applications such as quantum repeaters, where the memories facilitate long-range entanglement distribution, storage time becomes more critical given the time overheads in communication between nodes. The efficiency of the memory is also of critical importance, as this adds another loss channel which should be compared with direct transmission of the photon subject to loss. In the

simplest implementation, one might require lifetimes $\gtrsim 1$ s, and efficiencies $\gtrsim 90\%$ [76], however the lifetime threshold may be lowered if attempts to establish entanglement can be parallelised [38].

Furthermore, in order to reduce the time overheads needed to achieve heralded entanglement between remote QM nodes of the network, multimode capacity of the memory[20] becomes an appealing feature as proposed in [64], where if N photons can be stored in the memory then one may multiplex entanglement distribution attempts, offering an $N-$fold enhancement in success probability per experimental cycle.

Finally we mention some general properties of QMs. Firstly we must consider the fidelity of the state retrieved with the input (averaged over all possible input states) should surpass that of a reference process in the absence of the device. With the density matrices ρ_i describing the input/output states, the fidelity between the two states, defined as

$$F(\rho_i, \rho_j) = \mathrm{Tr}\left[\sqrt{\sqrt{\rho_i}\rho_j\sqrt{\rho_i}}\right]^2, \tag{1.1}$$

gives a monotonic metric of the similarity of the two states. This runs from zero if the states have no overlap, to being maximal at unity if the states are identical.

A suitable benchmark is to compare against a classical measure–and–prepare strategy, in which one measures the state directly, and on demand prepares a new state based upon the result of that measurement.[21] This gives a threshold fidelity $F > 2/3$ for a single photon qubit [78]. More stringent criteria are needed in order to preserve nonlocal correlations of a maximally entangled pair of qubits: this requires the fidelity to exceed $F \gtrsim 85\%$ in order to still violate the CHSH (Clauser-Horne-Shimony-Holt) inequality [79], where the CHSH inequality,[22] due to Clauser *et al.* [80], is an entanglement witness [81].

Meeting these minimum fidelity thresholds are not sufficient to guarantee outperforming direct transmission. Any noise introduced by the memory interaction (and from elsewhere in the system) must be eventually compensated for via an error correction scheme, or entanglement distillation protocol. This brings additional constraints on the minimum fidelity required such that any error introduced by the memory can be corrected for [82].

One useful metric to benchmark to compare memory protocols is the μ_1 [83] which is a simultaneous measure of noise and efficiency:

$$\mu_1 = \frac{\langle n_{\mathrm{noise}}\rangle}{\eta_{\mathrm{mem}}} \tag{1.2}$$

where η_{mem} is the total memory efficiency and $\langle n_{\mathrm{noise}}\rangle$ is the expectation of the number of noise photons accompanying the retrieved state. Equivalently the μ_1 expresses the input amplitude threshold resulting in a signal-to-noise ratio (SNR) of unity.

However, we note that the SNR is a "first order" measure and thus can't access the quantum statistics, and so metrics able to yield information about the photon correlations are also necessary; measures to provide distinction between different quantum and classical states of light include the cross-correlation $g^{(1,1)}$ and the auto-correlation $g^{(2)}$.

In particular we will return to the second-order autocorrelation function $g^{(2)}$ as a useful metric

[20] That is, a measure of how much quantum information a memory can store simultaneously [77], where each of those modes stored may be selectively addressed.

[21] In which case, the memory is purely classical.

[22] Originally proposed as a test that a bipartite system exhibits correlations which cannot be explained by a *local hidden-variable* model.

to distinguish different quantum and classical states, which can be expressed in terms of the photon ladder operators $\hat{a}^{(\dagger)}$ as

$$g^{(2)} = \frac{\langle \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} \rangle}{\langle \hat{a}^\dagger \hat{a} \rangle^2}. \tag{1.3}$$

If this is evaluated for a *classical field*[23] we find $g^{(2)} \geqslant 1$; a pure coherent state will have a $g^{(2)} = 1$, a thermal state $g^{(2)} = 2$, whilst a single photon Fock state $g^{(2)} = 0$. Measuring a $g^{(2)} < 1$ of the state retrieved from a memory thus verifies that the process has preserved the non-classical feature of the light see [e.g. 84, §IV.C].

In addition to the above considerations, in the interests of producing a more practical, scalable solution, one should also consider, for example, how such a memory might be manufactured ("can it be repeatably mass produced?"), and how is it operated ("what are the experimental overheads: do you need cryogenics, a vacuum system or a particularly bulky laser system?").

Finally, at some point in the process, the photons need to interface with a matter system, whether that be the system in which they are generated, the fibre through which they are transmitted, or some material node onto which they are mapped for subsequent operations: thus, depending on the application, there are preferred wavelengths we might wish to operate with, and hence want to be able to store in the memory.

1.4 Schemes & Platforms

Having established the usefulness of a quantum memory in principle, let us now consider QMs in practice. The simplest scheme, would be a delay line, whereby the light is stored in a switchable cavity; which would be inherently noise-free (due to the absence of atoms or other optical fields) and broadband since it doesn't rely upon any atomic resonance feature. The speed of the memory then depends on the bandwidth of the switch and the cavity roundtrip time – so one requires fast, low loss switches. This has been demonstrated in free space using a Pockels cell [85]; however the technical requirements to scale such a device are no small challenge.

The alternative is to store light in atomic, or atom-like media. This is possible in an isolated quantum system of a single emitter such as a trapped ion (e.g. [86]) or an atom in a cavity (e.g. [87]).

In many systems, however, in order to reach the strong light-matter coupling regime required for interaction with an isolated emitter, a high finesse cavity is needed. This then imposes a narrow acceptance bandwidth which in turn limits the achievable TBP. Perhaps then a more promising approach is to use an ensemble of emitters. In this case, the photon state is then stored by mapping it onto a collective excitation of the ensemble, and much higher values for the TBP can be attained. This is because, with a large ensemble, the absorption probability is increased without the need for many passes through a high finesse cavity. The acceptance bandwidth of the memory is then not limited by the linewidth of the cavity, so much shorter pulses in time can be stored.

Memories based on ensembles may be more technically simple to operate experimentally, since in this case there is no need to isolate a system of a single quantum emitter. Also, in some systems this may make the memory more resilient to noise or local inhomogeneities, for instance in off-resonant protocols (when $\Delta \gg \delta_{\text{IH}}$, where Δ is the off-resonant detuning and the δ_{IH} is the size of any inhomogeneous broadening). This makes the system more resilient to local variations of strain

[23] That is, a state $\rho_{\text{Cl.}}$ which may be expressed as a mixture of coherent states $|\alpha\rangle$ with a positive quasi-probability distribution $P(\alpha) \geqslant 0$ such that $\rho_{\text{Cl.}} = \int \mathrm{d}^2\alpha P(\alpha) |\alpha\rangle \langle\alpha|$.

or charge fluctuations which affect the energy structures of each emitter differently.

1.4.1 Storage Media

Alkali atoms are ideal as they provide an intrinsically "clean" system to work with. In order to isolate the atoms, these will be in a vapour cell at ambient temperatures, or held in an atomic trap (e.g. optical tweezers, magneto-optical trap, or lattice). Room temperature vapour cells are a technically simple system to interact with, however the atoms will be moving at thermal velocities and so the ensemble will exhibit a wide, Doppler broadened, inhomogeneous line width; whilst cooling and trapping atoms suppresses this thermal motion, it also brings with it additional experimental overheads (trapping lasers and vacuum systems).

In the solid state, ensembles such as Rare-Earth-ion (REI) doped crystals are an appealing medium since these offer very long lifetimes since the optically-addressable valence electrons are shielded from the environment by outer lying electron shells, and due to inhomogeneities in the crystal host typically exhibit broadenings of $\sim 0.1 \rightarrow 10$ GHz [88]. Defect centres in diamond, such as nitrogen-vacancies [89], have also received much interest as a promising platform for quantum memories. Solid state systems however, in general must be cryogenically cooled in order to suppress phonons. But also can offer the advantage of imposing selection rules due to crystal symmetry, which may be used to ensure only desired interaction occurs

1.4.2 Storage Schemes

Considering the requirements to store a photon, there are broadly two categories of operation principles for QMs in terms of the underlying physical mechanisms: (i) photon echo memories and (ii) adiabatic memories. These are introduced briefly in the following section; for further reading, see [84] for a recent review of optical QM protocols in atomic ensembles. A useful summary of different European efforts (in 2010) is given in [90], and [42] gives a review of applications of memories.

After outlining the broad families of operation principle, I show some of the current key results in the memories community tabulated in 1.1, and discuss where the present work sits within this landscape.

1.4.2(a) Photon Echo Protocols

Photon echo methods are the optical analogue of spin echo [91, 92]. This process requires a two level system, with inhomogenous broadening of the transition; where the broadening may be due to thermal motion of warm atoms, or the local strain environment of a solid state emitter. Light absorbed on a broadened transition will be stored as a superposition of coherences across the range of frequencies; these coherences precess at different rates, depending on their frequency and will lead to dephasing.

For a quantum memory, one seeks to control the rephasing of these coherences. This is generally achieved by externally controlling the broadening. Controlled reversible inhomogeneous broadening (CRIB) was first proposed in 2001 [93] whereby application of a spatially varying electric or magnetic fields broadens the transition inhomogeneously across the whole ensemble, where atoms at different positions experience a different shift in resonance. An incident signal resonant with the broadened transition will be mapped into a superposition of excitations across the ensemble – a collective spin wave. These excitation components will dephase according to their detuning. By changing the polarity of the external field at a time τ, the shifts are reversed and the coherences will rephase after another

time τ, and the signal re-emitted. Typically efficiencies are low ($\sim 0.1\%$) attributed to the difficulty in achieving high absorption over the broadened line [94], but also the reabsorption of the signal during retrieval.

The gradient echo memory (GEM) proposed in 2008 [95] is a close relative of CRIB in which the field gradient is applied parallel to the optical axis[24] – thereby suppressing re-absorption of the signal enabling near perfect memory efficiencies. A variant of this is $\Lambda-$GEM [96], which employs a third energy level for storage enabling longer storage times. The bandwidth of GEM schemes are determined by the gradient of the applied field and the length of the ensemble, and is typically on the order of ~ 1 MHz; however GEM memories are intrinsically highly multimode [97]. Although single photon storage has not yet been achieved in GEM, analysis of the quantum storage capability (in cold atoms) with coherent states [98] are promising.

Another approach within this family is the atomic frequency comb (AFC) [88]. In this case a frequency-periodic absorption structure is imprinted into the inhomogeneously broadened transition, via spectral hole burning or optical pumping. In this way once a signal is absorbed, the dipoles will naturally rephase, and the signal will be re-emited, after a time inversely related to the comb periodicity. Again AFC memories have high multimode capacity but typically operate on MHz bandwidth pulses. Of itself AFC simply constitutes a "delay line" since the retrieval time is set in advance, and so an additional step mapping the excitation to a storage state is required, for instance via rapid adiabatic passage (AFC+RAP), although this additional process typically limits the efficiencies which may be achieved.

1.4.2(b) Adiabatic Protocols

The adiabatic family of quantum memories require three level systems (typically in a Λ configuration), in which the light-matter interaction between the signal and the storage medium is mediated by an additional *classical* control field.

Electromagnetically induced transparency (EIT) was proposed for storage of light in 2000 [99]. A classical control field incident on the ensemble, near resonance with one arm of the Λ system, will dress the states, thereby creating a double resonance feature with a transparency window in between. A signal in resonance with the transparency window experiences a *reduction* in absorption.[25] By the Kramers-Kronig relations, this change in absorption is accompanied by a strong dispersive feature. This can be used as the basis for an optical memory, as this dispersion leads to reduction of the group velocity of the signal as it propagates through the ensemble, which slows and compresses the signal.[26] If the control field is adiabatically switched off whilst the signal is inside the medium, this collapses the EIT window – the dispersion diverges and the signal becomes stored in the medium as a collective atomic excitation.[27] Retrieval follows the time reverse of this process, and as the control is turned back on, the signal will be re-accelerated out of the storage medium. This interaction can be optimised to address a single mode via appropriate shaping of the control [100], and has been demonstrated [101].

Whilst there have been demonstrations of storage and retrieval for single photons using EIT in warm Rb [102], a separate work suggests that collisionally induced fluorescence from the excited state at longer storage times provides an unconditional noise floor which is too high for quantum

[24] Whereas CRIB employs transverse broadening.
[25] In principle, to zero.
[26] Such that a signal pulse longer than the medium may fit inside.
[27] A "dark state polariton", in this case.

applications, particularly problematic for thermal vapours [103]. Note that the acceptance bandwidth in EIT protocols will be set by the transparency window – if the signal is any wider than this it will suffer partial absorption and introduce a loss channel into the interaction – this sets a limit of typically ∼MHz.

This method has been extended to mapping the excitation to a Rydberg state (Ryd-EIT) [104, 105]. This is a promising route to implementing Rydberg-interaction mediated single-photon nonlinearities.

Much larger bandwidths can be achieved in an off-resonant Raman interaction between two metastable ground states [106]. Here the bandwidth is determined directly by the width of the control field, with an upper limit set by the ground-state splitting between the two ground states. To make the most of this, media with large ground state splittings are generally chosen, such as caesium (warm vapour) [107], phonons in bulk diamond [108] or molecular hydrogen [109]. Raman also offers the functionality of a multi-mode memory via wavevector multiplexing within the ensemble [110, 111].

Whilst on the topic of EIT and Raman interactions, some mention should be given to the Duan-Lukin-Cirac-Zoller (DLCZ) scheme [112], which might be described as a hybrid between a photon source and a memory. In this case a spin polarised Λ medium is pumped, until spontaneous Raman scattering heralds the creation of a long lived spin wave excitation, which can then be retrieved as a photon, on demand.[28] Such a scheme is an appealing one for the distribution of entanglement [113]. Whilst this doesn't allow for the storage of an *incident* signal from elsewhere, as a photon source it offers functionality such as the control over the spatio-temporal profile of the output photon retrieved from the heralded spin-wave [114].

A recent addition to the family are the off-resonant cascaded absorption (ORCA) [115] and/or fast ladder memory (FLAME) [116] protocols,[29] which are conceptually similar to the Raman scheme, but operate on a Ξ (ladder) energy scheme rather than Λ. This approach offers large bandwidths (storage of GHz bandwidth single photons has been demonstrated [115], but larger bandwidths are, in principle, possible [117]) and suffers no intrinsic noise processes, and has recently been demonstrated for the storage of single photons ($g^{(2)}_{\text{out}} \approx g^{(2)}_{\text{in}} = 0.02$ [115]); however, the current demonstrations are limited to lifetimes of ~ 10 ns in Cs [115], or ~ 86 ns in Rb [116]. The latter timescale here would be useful for local synchronisation applications.

1.4.2(c) Summary

Table 1.1 shows some of the recent benchmarks in performance, in various protocols, across the memories community.

Different memories[30] are naturally more suited to differing applications. For instance, EIT schemes typically offer very long lifetimes, whilst GEM protocols offer excellent efficiencies; both of these are generally narrow band whereas the Raman memory offers much higher bandwidths. In this thesis, we focus on the Raman protocol in warm caesium vapour.

We note that Λ schemes,[31] in isolation, are susceptible to four-wave mixing (FWM) noise – which for monatomic media prevents the faithful retrieval of single photons [e.g. 118, in the Raman protocol]; however, this may be avoided in solid state or molecular media in which additional selection rules arise and may be used to suppress unwanted interactions [108, 109], although these memories are limited

[28] And, for subsequent synchronisation, these photons are naturally compatible with a similar memory medium.
[29] These are similar protocols, differing only in detuning from the intermediate state, were both realised at the same time by different groups (in Rb [116] and in Cs [115]).
[30] By which I refer to both protocol and medium.
[31] GEM, EIT and Raman.

Protocol	Medium	Notes
GEM	Cold Rb	$\eta \approx 87\%$, $\tau \approx 1$ ms[98]
GEM	Warm Rb	$\mathcal{F} \approx 98\%$ for coh. input $\langle N_{\text{in}} \rangle = .67$ [122];
GEM	Warm Rb	$\eta \approx 87\%$ [121]
GEM	REI	$\eta \approx 69\%$ [123]
AFC	REI	quantum storage, $\eta \sim 20\%$ [124], quantum storage, 5 GHz, waveguide [125]
AFC+RAP	REI	quantum storage, $\mu_1 \approx 5\%$, $\tau \approx 30$ μs [126], $\eta \approx 5.7\%$ [127]
EIT	Warm Rb	quantum storage [102]
EIT	Cold Rb	$\tau \sim 1$ min, [128]
EIT	Cold Rb	quantum storage,$\tau \sim 11$ μs, $\eta \sim 6\%$, ~ 1 MHz [129]
EIT	REI	$\tau \sim$1 min [130],
Ryd-EIT	cold Rb	$\eta = 3.4\%$, $\tau = 3$ μs, 7 MHz, quantum storage,[105]
Ryd-EIT	cold Rb	$\eta = 23\%$, $\tau = 4$ μs, 10 MHz, quantum storage,[104]
Raman	RT Diamond	THz bandwidth, quantum storage [131]; spectral manipulation [119]
Raman	Warm Cs	$\eta \sim 30\%$, GHz bandwidth[118], $\tau \sim 150$ μs [107]
Raman	Warm H_2	$\eta \sim$10%, $\tau \sim$85 ps, THz bandwidth [109]
ORCA	Warm Cs	$\eta \sim 17\%$,$\tau \sim 5$ ns, $\mu_1 = 4 \times 10^{-5}$, ~ 1 GHz, quantum storage [115]
FLAME	Warm Rb	$\eta \sim$25%, $\tau \sim$86 ns, $\mu_1 = 2 \times 10^{-4}$, $\sim$600 MHz[116]

Table 1.1: Experimental quantum memories. Quoted values for: memory efficiency η, memory lifetime τ , frequencies give the acceptance bandwidth, and "quantum storage" indicates that *some* threshold of non-classicality has been surpassed – such as given in section IV.C of reference [84].

to sub-ns lifetimes. Addressing the FWM in alkali vapour is one of the themes in this thesis.

Multimode capacity is an oft cited merit of quantum memories as an additional degree of freedom for multiplexing; however, in this thesis we focus upon the potential applications of a *single mode* quantum memory, in which we may selectively address a particular state of light.

Furthermore, motivated by the long-term goal of realising a quantum network – in which it will be necessary to interface dissimilar devices – the ability to spectrally manipulate a quantum signal is a fundamental building block. This can be achieved with memories, and has been demonstrated in the Raman protocol in room temperature diamond (via phonons) [119] on THz bandwidth single photons showing bandwidth modulation by factors $0.5 \rightarrow 1.9$ and central frequency tunability of $\sim 4\times$ the input bandwidth. This is also possible in $\Lambda-$GEM [120] demonstrating a shift of 600 kHz on $\lesssim 1$ MHz bandwidth coherent states [121]. The Raman memory in warm caesium vapour also offers this functionality, but in a bandwidth range inaccessible by demonstrated platforms.

1.5 Thesis Outline

The results in this thesis are presented roughly in chronological order. In chapter 2, I introduce the Raman memory, and derive the equations of motion used in subsequent chapters.

In part I, I present my contributions to the theoretical description and first experimental demonstration of noise suppression in Raman memories (chapter 3), and some of my efforts to explore the optimisation of the cavity Raman memory (chapter 4).

Then, in part II, I move towards a discussion of temporal modes of light, within the wider landscape of different encodings for photonic quantum information in chapter 6; and finally, in chapter 7, I present our preliminary investigation of the modal structure of the memory interaction along with the tools developed, for the manipulation of temporal modes. Whilst not used directly in the main thesis, in appendix B I introduce a new technique I have developed for the characterisation of the temporal complex spectrum of pulses.

CHAPTER 2

THE RAMAN QUANTUM MEMORY

In this chapter I introduce the Raman memory protocol, and provide a brief derivation of the equations of motion; however, here I take full account of four-wave-mixing in the interaction, which must be considered as a noise source.

2.1 Introduction

By way of introduction, a simple picture of the Raman memory scheme used by this group is briefly outlined here, before we proceed to a more formal derivation in the next section.

The operation of an adiabatic Raman memory is based on a Λ energy level configuration, schematically shown in figure 2.1. The system is prepared in one of the hyperfine ground state levels $|1\rangle$ via optical pumping. Then to "write" a photon to the memory, a Raman transition between the ground state levels $|1\rangle \rightarrow |3\rangle$ detuned from resonance with the excited level $|2\rangle$, is driven using a strong control pulse, without the intermediate state being occupied, in a similar way to STIRAP [132].

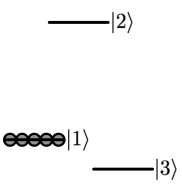
In this way, the optical signal, that is the photon state to be stored, is written into a long-lived ground state coherence, or spin-wave. The photon state is then "read" from the memory, on demand, by again applying another control pulse and driving the transition back to the original state with the emission of the signal photon. One notes that the storage and retrieval control pulses need not be identical – for instance by changing the detuning of the control between the storage and retrieval interaction also allows the user to modify the frequency of the retrieved signal, allowing the photon to be read out at a different colour [106, 119].

The selection rules for the Raman transition, far off resonance in monatomic vapour (in the absence of symmetry imposed restrictions) dictate that the two photons driving the transition have orthogonal polarisations [133], therefore the signal can straightforwardly be isolated from the control pulse after the memory medium using both polarization and frequency filtering.

An appropriate energy structure for this scheme is found in the ground state levels of atomic caesium, using the identification (e.g. refs [107, 134, 135]):

$$|1\rangle : 6^2S_{1/2}(F=4)$$
$$|2\rangle : 6^2P_{3/2}(\text{manifold})$$
$$|3\rangle : 6^2S_{1/2}(F=3)$$

$\longrightarrow$



where the hyperfine structure of $|2\rangle$ is not resolved in these experiments. In an ensemble of atomic caesium, storage of a photon is realized as a spin wave: an entangled superposition of states in which one[1] atom in the ensemble has its spin flipped.

Proof of principle experiments have demonstrated storage of light using this protocol both in a bulk cell [107], and in a hollow core photonic crystal fibre [135], showing a high TBP of 1000 required for a practical quantum memory; however current experiments are limited by the noise levels due to spontaneous four wave mixing.

2.1.1 Noise: Spontaneous Four-Wave Mixing

After a promising start for the Raman memory, an experiment published by our group in 2014 [118], just prior to my arrival, demonstrated the presence of a noise process inherent in the Raman memory scheme due to sponteneous four-wave mixing, the mechanics of which are discussed in detail in the following section and subsequent chapter.

The basic principle is summarized in the lower panel of figure 2.1 where even in the absence of an incident photon in the signal mode (blue), the control field (orange) may spontaneously scatter a photon into a different mode (green). This heralds the creation of a spurious spin excitation which may be retrieved into the signal mode. Overall then we have the spontaneous scattering of two control photons into two different modes – hence four-wave mixing (FWM) – in which one of the modes coincides with the signal mode. This results in unwanted noise in the signal mode emerging from the memory.

The four-wave-mixing satisfies energy and momentum conservation, and due to the difference in detuning from the atomic resonances, the three fields will experience different dispersion. Thus the Stokes/anti-Stokes photons will be emitted into slightly different spatial distributions [136] in a cone about the mode defined around the control field;[2] however this effect is small in our operating conditions [146].[3]

Notice that the initial configuration means that the detuning for the spurious process is *greater* than for the desired process, so the Raman coupling is suppressed with respect to the target memory write transition;[4] however, when this does occur and the stokes photon is emitted into the solid angle of detection, the resulting state is then indistinguishable from the system where a photon has been written normally to the memory. The output photon is in the same spatial, temporal and frequency mode as if a photon has been written to the memory via the intended transition. This gives rise to spurious counts even when no photon has been written to the memory, which can't be isolated from the successfully stored photons [135].

The noise from this process has been demonstrated to be prohibitively high for the storage of single photons [118], and furthermore, as we will show in the following sections, it is an intrinsic part of the memory Hamiltonian. Therefore, it must be eliminated during the memory interaction in order for the caesium ensemble to be a useful quantum memory.

[1] Or, at most one in the case of storage of a single photon – more if the input state contains more photons.

[2] Any seeding of the FWM by a single photon may be assumed negligible.

[3] An investigation into noise suppression by misaligning the signal and control; however, no suppression was seen before the memory efficiency was significantly degraded.

[4] This is also achievable if the ensemble is initialised into the lower ground state, and the memory transition is red detuned.

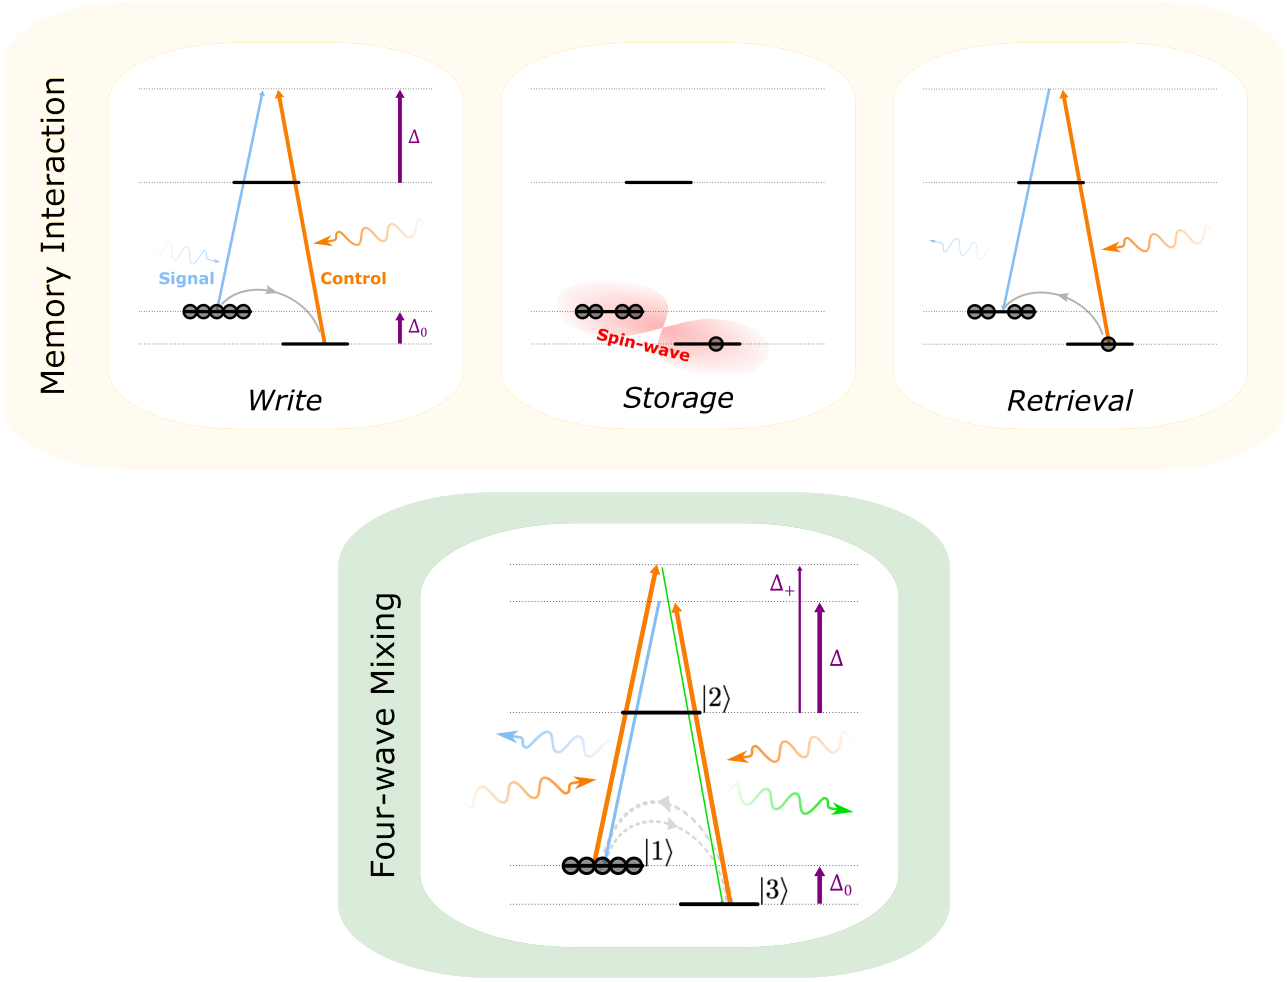


Figure 2.1: Top: Target memory interaction whereby the control field mediates a reversible mapping betwen the signal mode and spin-wave in the atomic ensemble. Bottom: Four-wave mixing process resulting in the production of a pair of photons, one of which is in the signal mode. Note that in our case we will consider the signal and control as *co-propagating*.

2.2 Equations of Motion

We now proceed to find a more formal description of the interactions occuring during the memory interaction, and in so doing try to capture the dynamics of both the storage interaction *and* the noise process. This derivation follows that presented in ref [133] for co-propagating signal and control, but modified here to include the possibility of four–wave mixing from the outset.

For this derivation, we consider an ensemble of atoms in free space with a Λ energy configuration, where the atomic eigenstates are labelled $|i\rangle$ for $i = 1, 2, 3$, and the atoms are optically pumped into the state $|1\rangle$, and we store the signal in an atomic coherence between the levels $|1\rangle$ and $|3\rangle$, where we show the target interaction in the upper panel of figure 2.1.

Firstly let us write down the optical fields present: we apply a strong control pulse at frequency ω_c which can be described by a classical field that is unaffected by the interaction with the atomic ensemble. The electric field vector $\boldsymbol{E}_\mathrm{c}$ is assumed to exhibit cylindrical symmetry as it propagates along the z axis as a function of time t, with a tranverse radial coordinate $\boldsymbol{\rho}$. The field may be written in terms of its envelope $\mathcal{E}_\mathrm{c}$, polarisation $\boldsymbol{v}_\mathrm{c}$, and carrier wave:

$$\boldsymbol{E}_\mathrm{c}(z,t,\boldsymbol{\rho}) = \mathcal{E}_\mathrm{c}(z,t,\boldsymbol{\rho})\boldsymbol{v}_\mathrm{c}\mathrm{e}^{\mathrm{i}\omega_c(t-z/c)} + \mathrm{c.c.} \tag{2.1}$$

Then, similarly, we need to define the weak signal field we wish to store: this must be treated with a quantum mechanical description since in general we aim to store nonclassical states of light such as a single photon, which can be written:

$$\boldsymbol{E}_{\mathrm{s}}(z,t,\boldsymbol{\rho}) = \mathrm{i}\boldsymbol{v}_{\mathrm{s}} \int \mathrm{d}\omega g(\omega)\hat{a}(\omega,t,\boldsymbol{\rho})\mathrm{e}^{-\mathrm{i}\omega_{\mathrm{s}}z/c} + \mathrm{h.c.} \tag{2.2}$$

where $\hat{a}^{(\dagger)}(\omega,t,\boldsymbol{\rho})$ are the Bosonic annilation (creation) field operators for a photon with frequency ω (at time t and radial coordinate $\boldsymbol{\rho}$). We omit the "hat" $\hat{\cdot}$ in the subsequent derivation for readability. Inside the integral $g(\omega) = \sqrt{\hbar\omega/4\pi\epsilon_0 c\mathcal{A}}$ is the mode density at frequency ω, where $\mathcal{A}$ is the beam area.

Provided that the bandwidth is much smaller than the central frequency ω_{s}, this can be rewritten in a slowly varying envelope picture as

$$\boldsymbol{E}_{\mathrm{s}} = \mathrm{i}\boldsymbol{v}_{\mathrm{s}} g_{\mathrm{s}} S(z,t,\boldsymbol{\rho})\mathrm{e}^{\mathrm{i}\omega_{\mathrm{s}}(t-z/c)} + \mathrm{h.c.} \tag{2.3}$$

where we approximate the field as monochromatic,[5] and is captured by the density of states $g_{\mathrm{s}} = \sqrt{2\pi}g(\omega_{\mathrm{s}})$ at the centre frequency, and the propagating field amplitude is written:

$$S(z,t,\boldsymbol{\rho}) = \mathrm{e}^{-i\omega_{\mathrm{s}}(t-z/c)} \times \frac{1}{\sqrt{2\pi}} \int \mathrm{d}\omega a(\omega,t,\boldsymbol{\rho})\mathrm{e}^{-\mathrm{i}\omega z/c}. \tag{2.4}$$

Having established our expression for the incident fields, we next consider the Hamiltonian for the total system, which may be expressed as:

$$\mathcal{H} = H_{\mathrm{L}} + H_{\mathrm{A}} + H_{\mathrm{int}}, \tag{2.5}$$

where we have separated out the Hamiltonian for the light field

$$H_{\mathrm{L}} = \int \mathrm{d}\omega \hbar\omega \hat{a}^{\dagger}(\omega)\hat{a}(\omega), \tag{2.6}$$

and for an atom within the atomic ensemble:

$$H_{\mathrm{A}} = \sum_{j=1}^{3} \hbar\omega_j \hat{\sigma}_{jj}, \tag{2.7}$$

where[6] $\sigma_{jk} = |j\rangle\langle k|$ and j, k label the atomic energy eigenstates. The diagonal elements σ_{jj} describe the atomic populations in each energy level, while the off-diagonal elements $\sigma_{j\neq k,k}$ capture the coherences between them.

Finally the interaction Hamiltonian which arises from the potential energy of charge displacement within an atom and the light field:

$$H_{\mathrm{int}} = -\boldsymbol{E}\cdot\boldsymbol{d}, \tag{2.8}$$

[5] By which I mean that the amplitude may be taken as constant across the signal spectrum.
[6] Again, we will omit the hats to avoid clutter...

where the dipole moment

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

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

where in the second line we have used that the dipole transition between the two ground states $|1\rangle$ and $|3\rangle$ is strongly forbidden.

Now for the dynamics. With this Hamiltonian, we wish to describe the time evolution of the signal field[7] and the atomic ensemble – we hence look at the time evolution of the atomic operators σ_{ij} using the relationship:

$$\partial_t \sigma_{ij} = \frac{\mathrm{i}}{\hbar}\left[\sigma_{ij}, \mathcal{H}\right], \tag{2.11}$$

Evaluating each of these terms ($j = 1, 2, 3$) yields a set of five independent equations of motion for the atomic operators, two for the populations (since the population is conserved and must satisfy $\sum_j \sigma_{jj} = \mathbb{1}$)

$$\partial_t \sigma_{11} = -\mathrm{i}\boldsymbol{E}\cdot\left[\boldsymbol{d}_{12}\sigma_{12} - \boldsymbol{d}_{12}^*\sigma_{21}\right] \tag{2.12}$$

$$\partial_t \sigma_{33} = +\mathrm{i}\boldsymbol{E}\cdot\left[\boldsymbol{d}_{23}\sigma_{23} - \boldsymbol{d}_{23}^*\sigma_{32}\right], \tag{2.13}$$

and three for the coherences:

$$\partial_t \sigma_{12} = \mathrm{i}\omega_{21}\sigma_{12} - \mathrm{i}\boldsymbol{E}\cdot\left[\boldsymbol{d}_{12}^*(\sigma_{11} - \sigma_{22}) + \boldsymbol{d}_{23}\sigma_{13}\right] \tag{2.14}$$

$$\partial_t \sigma_{13} = \mathrm{i}\omega_{31}\sigma_{13} - \mathrm{i}\boldsymbol{E}\cdot\left[\boldsymbol{d}_{23}^*\sigma_{12} - \boldsymbol{d}_{12}^*\sigma_{23}\right] \tag{2.15}$$

$$\partial_t \sigma_{23} = \mathrm{i}\omega_{32}\sigma_{23} - \mathrm{i}\boldsymbol{E}\cdot\left[\boldsymbol{d}_{23}^*(\sigma_{22} - \sigma_{33}) - \boldsymbol{d}_{12}\sigma_{13}\right], \tag{2.16}$$

in which we use the notation $\omega_{jk} = \omega_j - \omega_k$.

Before proceeding, note that through these coupled equations of motion, the control field may also couple the states $|1\rangle$ and $|2\rangle$ through $\boldsymbol{d}_{12}\cdot\boldsymbol{v}_\mathrm{c}$ and in so doing may generate coherence in σ_{32}. This is the origin of the four-wave mixing process which in previous derivations [133] was assumed to be negligible. So for book–keeping, we now allow an additional *weak* field in two–photon resonance with the control field in this configuration (i.e. at a detuning $\Delta_+ = \Delta + \Delta_0$, where Δ_0 is the ground state splitting – see table 2.1) which[8] like the signal field takes the form:

$$\boldsymbol{E}_\mathrm{a} = \mathrm{i}\boldsymbol{v}_\mathrm{a} g_\mathrm{a} A(z,t,\boldsymbol{\rho})\mathrm{e}^{\mathrm{i}\omega_\mathrm{a}(t-z/c)} + \text{h.c.} \tag{2.17}$$

where $\omega_\mathrm{a} = \omega_\mathrm{c} + (\omega_\mathrm{c} - \omega_\mathrm{s})$.

With these five coupled equations of motion, we can make a series of approximations [133] in order to find a more tractable and more physically insightful set of equations of motion:

1. **Linear approximation** (weak excitation):

Given that the stored signal is weak, the majority of the atoms remain in their initial state. We therefore replace the populations with the expectation values: $\sigma_{11} \to \nu \approx 1$, and $\sigma_{22} \to 0$ and $\sigma_{33} \to (1-\nu) \approx 0$,

[7] Recall, we assume the control field is strong enough such that it is unperturbed on passing through the system.

[8] Adopting the conventional terminology for Raman scattering, we name this the "anti–Stokes" field: in this configuration $\omega_\mathrm{s} < \omega_\mathrm{c}$, and the photon scattered in this process is of higher energy than the absorbed (control field) photon.

	signal	control	anti–Stokes
Δ_{+2}			$\omega_{21}-\omega_{\mathrm{a}}$
Δ_{+}		$\omega_{21}-\omega_{\mathrm{c}}$	$\omega_{23}-\omega_{\mathrm{a}}$
Δ	$\omega_{21}-\omega_{\mathrm{s}}$	$\omega_{23}-\omega_{\mathrm{c}}$	
Δ_{-}	$\omega_{23}-\omega_{\mathrm{s}}$		

Table 2.1: Labels for the respective detunings between the atomic and optical frequences.

where we have parametrised the pumping efficiency ν to allow for imperfect initialisation of the ensemble; and set their time derivatives to zero.

2. **Rotating wave approximation** (small detuning with respect to transition frequency):

We transform to the rotating frame and define $\tilde{\sigma}_{jk}=\sigma_{jk}\mathrm{e}^{\mathrm{i}\omega_{jk}\tau}$ where $\tau=t-z/c$ is the retarded time. Transforming as such,

$$\begin{aligned}\partial_t\sigma_{ij}&\to\partial_t\left(\mathrm{e}^{-\mathrm{i}\omega_{ij}\tau}\tilde{\sigma}_{ij}\right)\\&=-\mathrm{i}\omega_{ij}\mathrm{e}^{-\mathrm{i}\omega_{ij}\tau}\tilde{\sigma}_{ij}+\mathrm{e}^{-\mathrm{i}\omega_{ij}\tau}\partial_t\tilde{\sigma}_{ij},\end{aligned}$$

these two steps then yield

$$\partial_t\tilde{\sigma}_{12}=-\mathrm{i}\boldsymbol{E}\cdot\left[\nu\boldsymbol{d}_{12}^{*}\mathrm{e}^{-\mathrm{i}\omega_{21}\tau}+\boldsymbol{d}_{23}\tilde{\sigma}_{13}\mathrm{e}^{-\mathrm{i}\omega_{23}\tau}\right] \tag{2.18}$$

$$\partial_t\tilde{\sigma}_{13}=-\mathrm{i}\boldsymbol{E}\cdot\left[\boldsymbol{d}_{23}^{*}\tilde{\sigma}_{12}\mathrm{e}^{+\mathrm{i}\omega_{23}\tau}-\boldsymbol{d}_{12}^{*}\tilde{\sigma}_{23}\mathrm{e}^{-\mathrm{i}\omega_{21}\tau}\right] \tag{2.19}$$

$$\partial_t\tilde{\sigma}_{23}=-\mathrm{i}\boldsymbol{E}\cdot\left[(1-\nu)\boldsymbol{d}_{23}^{*}\mathrm{e}^{+\mathrm{i}\omega_{23}\tau}-\boldsymbol{d}_{12}\tilde{\sigma}_{13}\mathrm{e}^{+\mathrm{i}\omega_{21}\tau}\right]. \tag{2.20}$$

To complete the RWA, we then retain only terms oscillating slowly after expanding out the fields $\boldsymbol{E}$: i.e. terms with frequencies of the form $\pm(\omega_{A_i}-\omega_{L_j})$ are kept and terms of the form $\pm(\omega_{A_i}+\omega_{L_j})$ are taken to average to zero, where ω_{A_i} are the atomic resonance frequencies, and ω_{L_j} are the light field frequencies. We denote the respective detunings that remain as $\Delta_{\pm m}=\Delta\pm m\Delta_0$ where Δ is the two–photon detuning of the Raman transition, and Δ_0 is the ground state splitting.

Recalling the optical field definitions (the respective detunings are tabulated in 2.1), we are then left with:

$$\begin{aligned}\mathrm{i}\partial_t\tilde{\sigma}_{12}=&\left[\boldsymbol{s}\mathrm{e}^{-\mathrm{i}\Delta\tau}+\boldsymbol{c}\mathrm{e}^{-\mathrm{i}\Delta_{+}\tau}+\boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_{+2}\tau}\right]\cdot\boldsymbol{d}_{12}^{*}\nu\\&+\left[\boldsymbol{s}\mathrm{e}^{-\mathrm{i}\Delta_{-}\tau}+\boldsymbol{c}\mathrm{e}^{-\mathrm{i}\Delta\tau}+\boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_{+}\tau}\right]\cdot\boldsymbol{d}_{23}\tilde{\sigma}_{13}\end{aligned} \tag{2.21}$$

$$\begin{aligned}\mathrm{i}\partial_t\tilde{\sigma}_{13}=&\left[\boldsymbol{s}^{\dagger}\mathrm{e}^{+\mathrm{i}\Delta_{-}\tau}+\boldsymbol{c}^{*}\mathrm{e}^{+\mathrm{i}\Delta\tau}+\boldsymbol{a}^{\dagger}\mathrm{e}^{+\mathrm{i}\Delta_{+}\tau}\right]\cdot\boldsymbol{d}_{23}^{*}\tilde{\sigma}_{12}\\&-\left[\boldsymbol{s}\mathrm{e}^{-\mathrm{i}\Delta\tau}+\boldsymbol{c}\mathrm{e}^{-\mathrm{i}\Delta_{+}\tau}+\boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_{+2}\tau}\right]\cdot\boldsymbol{d}_{12}^{*}\tilde{\sigma}_{23}\end{aligned} \tag{2.22}$$

$$\begin{aligned}\mathrm{i}\partial_t\tilde{\sigma}_{23}=&\left[\boldsymbol{s}^{\dagger}\mathrm{e}^{+\mathrm{i}\Delta_{-}\tau}+\boldsymbol{c}^{*}\mathrm{e}^{+\mathrm{i}\Delta\tau}+\boldsymbol{a}^{\dagger}\mathrm{e}^{+\mathrm{i}\Delta_{+}\tau}\right]\cdot\boldsymbol{d}_{23}^{*}(1-\nu)\\&-\left[\boldsymbol{s}^{\dagger}\mathrm{e}^{+\mathrm{i}\Delta\tau}+\boldsymbol{c}^{*}\mathrm{e}^{+\mathrm{i}\Delta_{+}\tau}+\boldsymbol{a}^{\dagger}\mathrm{e}^{+\mathrm{i}\Delta_{+2}\tau}\right]\cdot\boldsymbol{d}_{12}\tilde{\sigma}_{13},\end{aligned} \tag{2.23}$$

where we have adopted the short–hand notation for the positive frequency component of the three field amplitudes (2.1, 2.3, 2.17) $\boldsymbol{c}=\mathcal{E}_{\mathrm{c}}(z,t,\boldsymbol{\rho})\boldsymbol{v}_{\mathrm{c}}$; and for the weak fields: $\boldsymbol{s}=\mathrm{i}g_{\mathrm{s}}S(z,t,\boldsymbol{\rho})\boldsymbol{v}_{\mathrm{s}}$, and $\boldsymbol{a}=\mathrm{i}g_{\mathrm{a}}A(z,t,\boldsymbol{\rho})\boldsymbol{v}_{\mathrm{a}}$.

Considering the physical interpretation of the atomic coherences, we first note from equation 2.10

that the polarisation in some volume element δV at position $\boldsymbol{r}$ is given by:

$$\begin{aligned}\boldsymbol{P} &= \frac{1}{\delta V}\sum_{\beta(\boldsymbol{r})}\boldsymbol{d}^{(\beta)} \\ &\approx \frac{1}{\delta V}\sum_{\beta(\boldsymbol{r})}\boldsymbol{d}_{12}\sigma_{12}^{(\beta)} + \text{h.c.}\end{aligned} \tag{2.24}$$

where β labels all of the atoms in δV at position $\boldsymbol{r}$, and we initially neglect the coherence between the two unoccupied states (and set $\sigma_{23} \approx 0$, but we will return to this in equation 2.28). We thus identify the *positive frequency* polarisation operator, driven at the signal frequency

$$\begin{aligned}P &= \frac{1}{\sqrt{n}\,\delta V}\sum_{\beta(\boldsymbol{r})}\sigma_{12}\mathrm{e}^{-\mathrm{i}\omega_\mathrm{s}\tau} \\ &= \frac{1}{\sqrt{n}\,\delta V}\sum_{\beta(\boldsymbol{r})}\tilde{\sigma}_{12}\mathrm{e}^{\mathrm{i}\Delta\tau},\end{aligned} \tag{2.25}$$

where n is the number density of atoms.

Similarly, we identify the spin wave coherence between the two dipole decoupled states, with an associated operator

$$B = \frac{1}{\sqrt{n}\,\delta V}\sum_{\beta(\boldsymbol{r})}\tilde{\sigma}_{13}. \tag{2.26}$$

With this definition, we note that both P and B satisfy bosonic commutation relations (at equal position/times) – if we consider the polarisation operator, for example:

$$\begin{aligned}\left[P(\boldsymbol{r},t),\, P^\dagger(\boldsymbol{r'},t')\right] &= \frac{1}{n\,(\delta V)^2}\sum_{\beta(\boldsymbol{r}),\gamma(\boldsymbol{r'})}\left[\tilde{\sigma}_{12}^{\beta}(t),\, \tilde{\sigma}_{21}^{\gamma}(t')\right] \\ &= \delta(t-t')\delta(|\boldsymbol{r}-\boldsymbol{r'}|)\times\frac{1}{\delta V}\times(\sigma_{11}-\sigma_{22}) \\ &= \delta(t-t')\delta(|\boldsymbol{r}-\boldsymbol{r'}|)\times\frac{1}{\delta V}\times 2\nu,\end{aligned} \tag{2.27}$$

and likewise for B. In the continuum limit $\delta V \to 0$ and with the linear approximation (replacing populations with expectation values) this reduces to $\left[P(\boldsymbol{r},t),\, P^\dagger(\boldsymbol{r'},t')\right] = \delta(t-t')\delta(|\boldsymbol{r}-\boldsymbol{r'}|)\times 2\nu$ and so we interpret these as annihilation operators for a collective atomic excitation at position $\boldsymbol{r}$ and time t.

We define an additional contribution to the polarisation – the remaining part of 2.10 – which is coupled to the additional, weak, anti–Stokes field, A:

$$R = \frac{1}{\sqrt{n}\,\delta V}\sum_{\beta(\boldsymbol{r})}\tilde{\sigma}_{32}\mathrm{e}^{\mathrm{i}\Delta_+\tau}, \tag{2.28}$$

which acts as a source term for A, i.e. a polarisation component oscillating at ω_a.

Combining the definitions for P (2.25) and B (2.26), with the equations of motion 2.21 – 2.23, we

then obtain:[9]

$$\begin{aligned}
\partial_t P(\boldsymbol{r},\tau) &= +\mathrm{i}\Delta P + \frac{1}{\sqrt{n}\delta V}\sum \mathrm{e}^{\mathrm{i}\Delta\tau}\partial_t\tilde{\sigma}_{12} \\
&= +\mathrm{i}\Delta P - \mathrm{i}\left[\boldsymbol{s} + \boldsymbol{c}\mathrm{e}^{-\mathrm{i}\Delta_0\tau} + \boldsymbol{a}\mathrm{e}^{-2\mathrm{i}\Delta_0\tau}\right]\cdot\boldsymbol{d}_{12}^*\nu - \mathrm{i}\left[\boldsymbol{s}\mathrm{e}^{+\mathrm{i}\Delta_0\tau} + \boldsymbol{c} + \boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_0\tau}\right]\cdot\boldsymbol{d}_{23}B \\
\partial_t R(\boldsymbol{r},\tau) &= +\mathrm{i}\Delta_+ R + \frac{1}{\sqrt{n}\delta V}\sum \mathrm{e}^{\mathrm{i}\Delta_+\tau}\left(\partial_t\tilde{\sigma}_{23}\right)^\dagger \\
&= +\mathrm{i}\Delta_+ R + \mathrm{i}\left[\boldsymbol{s}\mathrm{e}^{+2\mathrm{i}\Delta_0\tau} + \boldsymbol{c}\mathrm{e}^{+\mathrm{i}\Delta_0\tau} + \boldsymbol{a}\right]\cdot\boldsymbol{d}_{23}(1-\nu) - \mathrm{i}\left[\boldsymbol{s}\mathrm{e}^{+\mathrm{i}\Delta_0\tau} + \boldsymbol{c} + \boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_0\tau}\right]\cdot\boldsymbol{d}_{12}^* B^\dagger \\
\partial_t B(\boldsymbol{r},\tau) &= \frac{1}{\sqrt{n}\delta V}\sum \partial_t\tilde{\sigma}_{13} \\
&= -\mathrm{i}\left[\boldsymbol{s}^\dagger\mathrm{e}^{-\mathrm{i}\Delta_0\tau} + \boldsymbol{c}^* + \boldsymbol{a}^\dagger\mathrm{e}^{+\mathrm{i}\Delta_0\tau}\right]\cdot\boldsymbol{d}_{23}^* P + \mathrm{i}\left[\boldsymbol{s}\mathrm{e}^{+\mathrm{i}\Delta_0\tau} + \boldsymbol{c} + \boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_0\tau}\right]\cdot\boldsymbol{d}_{12}^* R^\dagger.
\end{aligned} \tag{2.29}$$

We can now manually include terms for decay to the equations of motion [133]: the excited state $|2\rangle$ has a natural linewidth, and its population decays at a rate 2γ, hence the coherences P and R decay as γ. So in the above, we replace the respective detunings $-\mathrm{i}\Delta_j \to \Gamma_j = \gamma - \mathrm{i}\Delta_j$ with a complex detuning for each field. In practice B also decoheres, but is neglected here as by assumption this is negligible on the timescale of the storage process.

In handling the decay in this fashion, we should incorporate Langevin noise operators. These ensure that the fluctuations responsible for this decoherence preserve the bosonic commutation relations of P, B and R as they couple to the electromagnetic field in spontaneously emitting. However, these are initially in the vacuum state, and the expectation values of these terms evaluate to zero [133] and may be neglected.

We now make the approximation that in the adiabatic $\Delta \gg \Omega$, δ_c and Raman limits $\Delta \gg \gamma$, the system adiabatically follows the applied fields and we can eliminate the excited state. This simply means that the detuning is the characteristic frequency and defines the dynamics: in this limit we can consider that Ω changes sufficiently slowly such that the excited state is never populated and only take explicit time–dependence from the exponentials when integrating equations 2.29:

$$P \approx \mathrm{i}\left\{-\nu\boldsymbol{d}_{12}\cdot\left[\frac{\boldsymbol{s}}{\Gamma} + \frac{\boldsymbol{c}\mathrm{e}^{-\mathrm{i}\Delta_0\tau}}{\Gamma_+} + \frac{\boldsymbol{a}\mathrm{e}^{-2\mathrm{i}\Delta_0\tau}}{\Gamma_{+2}}\right] - B\boldsymbol{d}_{23}\cdot\left[\frac{\boldsymbol{s}\mathrm{e}^{+\mathrm{i}\Delta_0\tau}}{\Gamma_-} + \frac{\boldsymbol{c}}{\Gamma} + \frac{\boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_0\tau}}{\Gamma_+}\right]\right\}, \tag{2.30}$$

$$R \approx \mathrm{i}\left\{(1-\nu)\boldsymbol{d}_{23}\cdot\left[\frac{\boldsymbol{s}\mathrm{e}^{+2\mathrm{i}\Delta_0\tau}}{\Gamma_-} + \frac{\boldsymbol{c}\mathrm{e}^{+\mathrm{i}\Delta_0\tau}}{\Gamma} + \frac{\boldsymbol{a}}{\Gamma_+}\right] - B^\dagger\boldsymbol{d}_{12}^*\cdot\left[\frac{\boldsymbol{s}\mathrm{e}^{+\mathrm{i}\Delta_0\tau}}{\Gamma} + \frac{\boldsymbol{c}}{\Gamma_+} + \frac{\boldsymbol{a}\mathrm{e}^{-\mathrm{i}\Delta_0\tau}}{\Gamma_{+2}}\right]\right\}. \tag{2.31}$$

Substituting these into 2.29, we find the time evolution of the spin–wave coherence:

$$\begin{aligned}
\partial_t B &= \left[\boldsymbol{c}^*\cdot\boldsymbol{d}_{23}^*\left(\frac{-\nu\,\boldsymbol{d}_{12}}{\Gamma} + \frac{(1-\nu)\,\boldsymbol{d}_{12}^*}{\Gamma^*}\right)\cdot\boldsymbol{s} + \boldsymbol{c}\cdot\left(\frac{-\nu\boldsymbol{d}_{12}}{\Gamma_+} + \frac{(1-\nu)\,\boldsymbol{d}_{12}^*}{\Gamma_+^*}\right)\boldsymbol{d}_{23}^*\cdot\boldsymbol{a}^\dagger\right] \\
&\quad - \left[\frac{|\boldsymbol{c}\cdot\boldsymbol{d}_{23}|^2}{\Gamma} + \frac{|\boldsymbol{c}\cdot\boldsymbol{d}_{12}^*|^2}{\Gamma_+^*}\right]B \\
&= \nu\left[\frac{-\mathrm{i}\Omega^{(23)\,*}\chi_{\mathrm{s}}^{(12)}}{\Gamma_{\mathrm{s}}}S + \frac{+\mathrm{i}\Omega^{(12)}\chi_{\mathrm{a}}^{(23)\,*}}{\Gamma_{\mathrm{a}}}A^\dagger\right] + \mathcal{O}(1-\nu) - \left[\frac{|\Omega^{(23)}|^2}{\Gamma_{\mathrm{s}}} + \frac{|\Omega^{(21)}|^2}{\Gamma_{\mathrm{a}}^*}\right]B,
\end{aligned} \tag{2.32}$$

where again, in the rotating wave approximation, we have neglected the oscillating terms – alternatively, this can be seen by discarding terms second order in the weak fields $\{S, B, P\}$. $\Delta_{(+)}$ is the two–photon detuning from resonance for the target (spurious) transitions; Ω_i are the Rabi frequencies

[9] For the sake of compactness and readability, I only explicitly include the time (t) and spatial ($\boldsymbol{r}$) dependence of the fields on the left side of the equations.

of the control field coupled to each of the transitions:

$$\Omega^{(ij)} = \frac{\boldsymbol{d}_{ij}.\boldsymbol{v}_c\,\mathcal{E}_c}{\hbar}, \tag{2.33}$$

where $\chi_{\mathrm{s,a}}^{(ij)}$ are the coupling constants for the interaction:

$$\chi_{\mathrm{s,a}}^{(ij)} = \left(\frac{\boldsymbol{d}_{ij}.\boldsymbol{v}_{\mathrm{s,a}}}{\hbar}\right) g_{\mathrm{s,a}}\sqrt{n}. \tag{2.34}$$

3. **Linear approximation 2**:

   Furthermore, in this perturbative limit, the coherences $\sigma_{j\neq k,k}$ and signal field S are small so after expanding out the field $\boldsymbol{E} = \boldsymbol{E}_{\mathrm{c}} + \boldsymbol{E}_{\mathrm{s}}$, we may then neglect second order terms (of the form $\sigma_{j\neq k,k}S$).

In addition to this, we may consider the propagation of the weak fields – taking the signal as an example – which from Maxwell's equation satisfies

$$\left(\nabla^2 - \frac{1}{c^2}\partial_t^2\right)\boldsymbol{E}_{\mathrm{s}} = \mu_0\partial_t^2\boldsymbol{P}_{\mathrm{s}}, \tag{2.35}$$

where $\boldsymbol{P}_s$ is the component of atomic polarisation which acts as a source term for the signal field (i.e. oscillating at frequency ω_s). We now make two further approximations to enable us to handle the propagation of the optical fields:

4. **"slow varying envelope" approximation**:

   This allows us to treat the propagation of the signal purely in terms of the envelope function, without considering the fast oscillating carrier wave (valid if the signal pulse duration is much greater than the optical period $2\pi/\omega_{\mathrm{s}}$).

We therefore express the "slowly varying" atomic polarisation – denoted by $\tilde{\cdot}$, factoring out the oscillation at the signal frequency (taking only take the positive frequency component of the macroscopic polarisation $\boldsymbol{P}$):

$$\boldsymbol{P}_{\mathrm{s}} = \tilde{\boldsymbol{P}}_{\mathrm{s}}\mathrm{e}^{\mathrm{i}\omega_{\mathrm{s}}\tau}, \tag{2.36}$$

where we note that with this definition, $\tilde{\boldsymbol{P}}_{\mathrm{s}}$ is related to the polarisation operators P and R through:

$$\begin{aligned}\tilde{\boldsymbol{P}}_{\mathrm{s}} &= \mathrm{e}^{-\mathrm{i}\omega_{\mathrm{s}}\tau} \times \frac{1}{\delta V}\sum_{\beta(\boldsymbol{r})}\left[\boldsymbol{d}_{12}\sigma_{12}^{(\beta)} + \boldsymbol{d}_{32}\sigma_{32}^{(\beta)}\right] \\ &= \sqrt{n}\left[\boldsymbol{d}_{12}P + \mathrm{e}^{-2\mathrm{i}\Delta_0\tau}\boldsymbol{d}_{32}R\right].\end{aligned} \tag{2.37}$$

If we then write down the propagation of the slowly varying field

$$\left(\nabla^2 - \frac{1}{c^2}\partial_t^2\right)\left[S\mathrm{e}^{\mathrm{i}\omega_{\mathrm{s}}\tau}\right]\boldsymbol{v}_{\mathrm{s}} \propto \partial_t^2\left[\tilde{\boldsymbol{P}}_{\mathrm{s}}\mathrm{e}^{\mathrm{i}\omega_{\mathrm{s}}\tau}\right], \tag{2.38}$$

where, using the chain rule, and applying the SVE approximation we can express the left–hand side

as

$$\left(\nabla^2 - \frac{1}{c^2}\partial_t^2\right)\left[S\mathrm{e}^{\mathrm{i}\omega_\mathrm{s}\tau}\right] = \mathrm{e}^{\mathrm{i}\omega_\mathrm{s}\tau}\left\{\nabla_\perp^2 + \left(\partial_z^2 - \frac{1}{c^2}\partial_t^2\right) - \frac{2\mathrm{i}\omega_\mathrm{s}}{c}\left(\partial_z + \frac{1}{c}\partial_t\right)\right\} S$$
$$\approx \mathrm{e}^{\mathrm{i}\omega_\mathrm{s}\tau}\left\{\nabla_\perp^2 - 2\mathrm{i}k_\mathrm{s}\left(\partial_z + \frac{1}{c}\partial_t\right)\right\} S, \tag{2.39}$$

and, likewise, the right:

$$\partial_t^2\left[\tilde{\boldsymbol{P}}_\mathrm{s}\mathrm{e}^{\mathrm{i}\omega_\mathrm{s}\tau}\right] = \mathrm{e}^{\mathrm{i}\omega_\mathrm{s}\tau}\left\{\partial_t^2 + 2\mathrm{i}\omega_\mathrm{s}\partial_t - \omega_\mathrm{s}^2\right\}\tilde{\boldsymbol{P}}_\mathrm{s}$$
$$\approx -\mathrm{e}^{\mathrm{i}\omega_\mathrm{s}\tau}\omega_\mathrm{s}^2\tilde{\boldsymbol{P}}_\mathrm{s}. \tag{2.40}$$

With these, after gathering together the coefficients, we then arrive at

$$\left\{\frac{\mathrm{i}}{2k_\mathrm{s}}\nabla_\perp^2 + \left(\partial_z + \frac{1}{c}\partial_t\right)\right\} S = -\left(\frac{\mu_0\omega_\mathrm{s}^2}{2g_s k_s}\right)\boldsymbol{v}_\mathrm{s}^* \cdot \tilde{\boldsymbol{P}}_\mathrm{s}, \tag{2.41}$$

and substituting in the expression for $\tilde{\boldsymbol{P}}_\mathrm{s}$, we find:

$$\left\{\frac{\mathrm{i}}{2k_\mathrm{s}}\nabla_\perp^2 + \left(\partial_z + \frac{1}{c}\partial_t\right)\right\} S = -\mathcal{A}\left[\chi_\mathrm{s}^{(21)*}P(z,\tau) + \chi_\mathrm{s}^{(23)*}R(z,\tau)\mathrm{e}^{-2\mathrm{i}\Delta_0\tau}\right]. \tag{2.42}$$

Now, we recall that the optical field $S = S(z,t,\boldsymbol{\rho})$ and has a transverse spatially varying profile, we make a final approximation:

5. **Paraxial approximation**:

   For this we assume that the beam propagates along z with negligible divergence and adopt a one dimensional description (valid if the transverse spatial profile $\mathcal{A}$ is much greater than the signal wavelength $2\pi/k_\mathrm{s}$), we may integrate over the transverse coordinate $\boldsymbol{\rho}$ and the transverse gradient, which describes the diffraction of the beam, approximates to zero.

Applying this to all of the fields,[10] and putting this all together, and likewise for the anti–Stokes, we arrive at two propagation equations for the weak optical fields

$$\left[\partial_z + \frac{1}{c}\partial_t\right] S(z,\tau) = -\chi_\mathrm{s}^{(21)*}P(z,\tau) - \chi_\mathrm{s}^{(23)*}R(z,\tau)\mathrm{e}^{-2\mathrm{i}\Delta_0\tau} \tag{2.43}$$
$$\left[\partial_z + \frac{1}{c}\partial_t\right] A(z,\tau) = -\chi_\mathrm{a}^{(23)*}R(z,\tau) - \chi_\mathrm{a}^{(21)*}P(z,\tau)\mathrm{e}^{+2\mathrm{i}\Delta_0\tau}. \tag{2.44}$$

Inserting in our expressions 2.30 and 2.31 (retaining only the non-rotating terms) into 2.43, and quoting again 2.29, we have:

$$\left[\partial_z + \frac{1}{c}\partial_t\right] S(z,\tau) = \left[-\nu\frac{\chi_\mathrm{s}^{(21)*}\chi_\mathrm{s}^{(12)}}{\Gamma_\mathrm{s}} + \mathcal{O}(1-\nu)\right] S + \mathrm{i}\frac{\chi_\mathrm{s}^{(21)*}\Omega^{(23)}}{\Gamma_\mathrm{s}}B \tag{2.45}$$
$$\left[\partial_z + \frac{1}{c}\partial_t\right] A(z,\tau) = \left[-\nu\frac{\chi_\mathrm{a}^{(21)*}\chi_\mathrm{a}^{(12)}}{\Gamma_\mathrm{a}^+} + \mathcal{O}(1-\nu)\right] A + \mathrm{i}\frac{\chi_\mathrm{a}^{(21)*}\Omega^{(21)}}{\Gamma_\mathrm{a}}B^\dagger, \tag{2.46}$$

and can write down our final equations of motion (setting $\nu = 1$) as:

[10] The fields A, S and B, as well as the classical field, and the number density n, is now a longitudinal density.

$$\begin{aligned}\left[\partial_z + \frac{1}{c}\partial_t\right] S(z,\tau) &= \frac{-\kappa_\mathrm{s}}{c} S + \mathrm{i}\frac{\chi_\mathrm{s}^{(21)*}\Omega^{(23)}}{\Gamma_\mathrm{s}} B \\ \left[\partial_z + \frac{1}{c}\partial_t\right] A(z,\tau) &= \frac{-\kappa_\mathrm{a}}{c} A + \mathrm{i}\frac{\chi_\mathrm{a}^{(21)*}\Omega^{(21)}}{\Gamma_\mathrm{a}} B^\dagger \\ \partial_t B(z,\tau) &= -\mathrm{i}\left[\frac{\chi_\mathrm{s}^{(12)}\Omega^{(23)*}}{\Gamma_\mathrm{s}} S - \frac{\chi_\mathrm{a}^{(23)*}\Omega^{(12)}}{\Gamma_\mathrm{a}} A^\dagger\right] - \left[\frac{\left|\Omega^{(23)}\right|^2}{\Gamma_\mathrm{s}} + \frac{\left|\Omega^{(21)}\right|^2}{\Gamma_\mathrm{a}^*}\right] B\end{aligned} \tag{2.47}$$

in which, we have defined

$$\kappa_\mathrm{s} = \frac{\chi_\mathrm{s}^{(21)*}\chi_\mathrm{s}^{(12)}}{\Gamma_\mathrm{s}} \tag{2.48}$$

$$\kappa_\mathrm{s} = \frac{\chi_\mathrm{a}^{(21)*}\chi_\mathrm{a}^{(12)}}{\Gamma_\mathrm{a}^+}. \tag{2.49}$$

In the subsequent analysis, and following chapter, for the purposes of discussing the memory dynamics we generally assume ideal pumping[11] and set $\nu = 1$.

2.3 Ideal Operation

To illustrate the physical characteristics of the Raman memory, for the remainder of this chapter we consider an idealised case where the spurious interactions may be neglected and we set $A = 0$, and drop the subscripts, s, for the signal field parameters. This leaves us with[12]

$$\begin{aligned}\partial_z S(z,\tau) &= \frac{-\chi^2}{c\Gamma} S + \frac{\mathrm{i}\chi\Omega}{\Gamma} B \\ \partial_\tau B(z,\tau) &= \frac{-\mathrm{i}\chi\Omega^*}{\Gamma} S + \frac{-|\Omega|^2}{\Gamma} B,\end{aligned} \tag{2.50}$$

Where, at the same time, we have transformed to work in terms of the retarded time $\tau = t - z/c$, such that

$$\partial_z + \frac{1}{c}\partial_t = \partial_z|_\tau\,, \qquad \partial_t = \partial_\tau|_z\,. \tag{2.51}$$

Furthermore, we move to dimensionless units denoted with a tilde, where temporal/spatial variables are scaled with respect to the natural linewidth γ and the interaction length L. Since we consider the field modes in terms of the boundary conditions: the signal field in terms of its temporal profile $S_\mathrm{in}(\tau) = S_\mathrm{in}(z=0,\tau)$ and the spin wave in terms of its spatial profile $B_\mathrm{mem}(z) = B_\mathrm{mem}(z, \tau\to\infty)$. The normalisation then must preserve terms of the form $\int_{-\infty}^{\infty}\mathrm{d}\tau S_\mathrm{in}^\dagger(\tau)S_\mathrm{in}(\tau)$, and likewise for B, which are dimensionless and just yield the number of excitations.

So, as a result of transformation when we change to dimensionless units for time: $\tilde{\tau} = \gamma\tau$ we can see that

$$\tilde{S} = S/\sqrt{\gamma}, \quad \tilde{\Omega} = \Omega/\gamma, \quad \tilde{\Gamma} = \Gamma/\gamma. \tag{2.52}$$

[11] This will be motivated in practice in section 3.4.4, from which we conclude that FWM is the dominant noise process. And, moving forwards in section 4.3 we confirm pumping efficiencies can be achieved in excess of 99%.

[12] Also, we assume χ is real, WLOG – and the superscript index ordering becomes unimportant.

By the same argument taking $\tilde{z} = z/L$ then

$$\tilde{B} = B\sqrt{L}, \tag{2.53}$$

and finally

$$\tilde{\chi} = \chi\sqrt{\frac{L}{\gamma}} = \sqrt{d}, \tag{2.54}$$

which defines the optical depth of the ensemble.

The final equations of motion, in normalized units then become

$$\left(\partial_{\tilde{z}} + \frac{d}{\tilde{\Gamma}}\right)\tilde{S} = +\mathrm{i}\frac{\tilde{\Omega}\sqrt{d}}{\tilde{\Gamma}}\tilde{B} \tag{2.55}$$

$$\left(\partial_{\tilde{\tau}} + \frac{|\tilde{\Omega}|^2}{\tilde{\Gamma}}\right)\tilde{B} = -\mathrm{i}\frac{\tilde{\Omega}^*\sqrt{d}}{\tilde{\Gamma}}\tilde{S}. \tag{2.56}$$

With these, since the memory protocol is linear, we can then characterise the mapping from one mode into the other via a set of Green's functions:

$$B_{\mathrm{mem}}(z) = \int_{-\infty}^{\infty} \mathrm{d}\tau K_1(z,\tau) S_{\mathrm{in}}(\tau) \tag{2.57}$$

$$S_{\mathrm{ret}}(\tau) = \int_0^1 \mathrm{d}z K_2(z,\tau) B_{\mathrm{mem}}(z), \tag{2.58}$$

where I label S_{in} for the input field, B_{mem} the spin-wave excited in the atomic ensemble and S_{ret} the retrieved optical field. Overall, the retrieved signal may be written:

$$S_{\mathrm{ret}}(\tau) = \int_{-\infty}^{\infty} \mathrm{d}\tau' K_{\mathrm{total}}(\tau,\tau') S_{\mathrm{in}}(\tau'), \tag{2.59}$$

where the Green's functions encode the details of the control pulse. The efficiency is optimised by the overlap between S_{in} and the control pulse. In this idealised case, with the respective limits satisfied, the Green's functions may be analytically obtained, which we will make use of later.

With this picture, we can make a useful definition for the memory efficiency in terms of the number of excitations:

$$N_{\mathrm{in}} = \int_{-\infty}^{\infty} \mathrm{d}\tau\, S_{\mathrm{in}}^{\dagger}(\tau) S_{\mathrm{in}}(\tau) \tag{2.60}$$

$$N_B = \int_0^1 \mathrm{d}z\, B_{\mathrm{mem}}^{\dagger}(z) B_{\mathrm{mem}}(z), \tag{2.61}$$

with which we can then consider the efficiencies for the various stages of the protocol:

$$\begin{aligned} \text{storage efficiency:} \quad & \eta_{\mathrm{store}} &&= N_B/N_{\mathrm{in}} \\ \text{retrieval efficiency:} \quad & \eta_{\mathrm{ret}} &&= N_{\mathrm{out}}/N_B \\ \text{total memory efficiency:} \quad & \eta_{\mathrm{total}} &&= N_{\mathrm{out}}/N_{\mathrm{in}} \quad = \quad \eta_{\mathrm{store}}\,\eta_{\mathrm{ret}}. \end{aligned}$$

2.3.1 Memory

In the adiabatic limit, the equations of motion reduce to a "beam splitter" type relation between the propagating signal mode and the stationary spin wave mode of the atomic ensemble.

This analogy can be seen by expressing the equations of motion for the operators for the electric field $\tilde{S}(\tilde{z},\tilde{\tau})$ and spin wave $\tilde{B}(\tilde{z},\tilde{\tau})$ in terms of the unitary[13] transformation [133]:

$$\tilde{S}(\tilde{z},\tilde{\tau}) = \frac{\tilde{\Omega}(\tilde{\tau})}{\sqrt{\tilde{W}}}\sigma(\tilde{z},\omega)\mathrm{e}^{-\left(\frac{\tilde{W}\omega+d\tilde{z}}{\tilde{\Gamma}}\right)} \tag{2.62}$$

$$\tilde{B}(\tilde{z},\tilde{\tau}) = \beta(\tilde{z},\omega)\mathrm{e}^{-\left(\frac{\tilde{W}\omega+d\tilde{z}}{\tilde{\Gamma}}\right)}, \tag{2.63}$$

where W is the normalization $\tilde{W} = \int_{-\infty}^{\infty}\mathrm{d}\tilde{\tau}|\tilde{\Omega}(\tilde{\tau})|^2 = W/\gamma$, and $\omega(\tilde{\tau}) = (1/\tilde{W})\int_{-\infty}^{\tilde{\tau}}\mathrm{d}\tilde{\tau}'|\tilde{\Omega}(\tilde{\tau}')|^2$.

These then yield the equations of motion

$$\partial_{\tilde{z}}\sigma = -\tilde{C}\beta \tag{2.64}$$

$$\partial_{\omega}\beta = +\tilde{C}\sigma, \tag{2.65}$$

where C is the Raman coupling constant: $\tilde{C} = \sqrt{\tilde{W}d}/\tilde{\Gamma}$ in dimensionless units, or $C = \sqrt{Wd\gamma}/\Gamma$ in "normal units", where the optical depth d is defined as in equations (2.55) and (2.56).

In practice, d is simply related to the number of atoms with which the light field can interact, defined in terms of the on resonance absorption. For efficient light matter interaction we require the Raman coupling coefficient of the order unity.

These equations 2.64–2.65, are analogous to the the Heisenberg equations of motion one derives from the beam-splitter Hamiltonian: $\mathcal{B} = i\hbar\zeta a^{\dagger}b + \text{h.c.}$ [137] acting on modes a and b, with a coupling strength ζ (with reflectivity and transmission $\sim \sin[\zeta]$ and $\sim \cos[\zeta]$ respectively). In this case the partial derivatives now correspond to different coordinates. This means that we are instead mapping between a propagating signal mode into a stationary spin wave. This action hence defines a memory, where the coupling between the two modes is mediated by the control pulse, with an efficiency (c.f. beam-splitter reflectivity) determined by the Raman coupling constant.

It can be shown [100], that for this system, in the adiabatic limit ($\Delta \gg \Omega$) the maximimum achievable efficiency is determined solely by the optical depth of the medium,[14] independent of detuning, for any (smooth) input mode A_{in} provided one appropriately shapes the control field. The actual efficiency, however, depends on the details of the control and signal pulse, such as their spatial and temporal overlap.

Without going into the more rigorous analysis given by Gorshkov *et al.* in the series of papers [138–141], this can be argued in simple terms by considering the Raman coupling constant

$$|C|^2 = \frac{Wd\gamma}{|\Gamma|^2}, \tag{2.66}$$

where we assume

$$\begin{aligned} W &= \int_{-\infty}^{\infty}\mathrm{d}\tau|\Omega(\tau)|^2 \\ &\approx \left(|\Omega_{\text{max}}|^2/\delta_c\right), \end{aligned} \tag{2.67}$$

where we have taken the control to be a top-hat function with δ_c the bandwidth of the control pulse. Then noting that we are operating in the adiabatic limit we have $|\Gamma| \approx |\Delta|$.

[13] In the limit of large detuning, Δ.

[14] or in the presence of a cavity this is enhanced by the cavity finesse, as will be discussed in the next chapter.

Together these give

$$|C|^2 \approx \left(\frac{d\gamma}{\delta_c}\right)\left|\frac{\Omega_{\max}}{\Delta}\right|^2, \tag{2.68}$$

and, since we have $\Omega/\Delta \ll 1$, for efficient memory operation requiring $C^2 \gtrsim 1$ we therefore need $d\gamma/\delta_c \gg 1$. In our experiments, e.g. [118, 142], we typically operate with a bandwidth of $\delta_c = 1.2\,\text{GHz}$ at a detuning of $\Delta = 15.2\,\text{GHz}$, and so this condition corresponds to requiring an optical depth $d \gg\sim 100$ in order to achieve $C^2 \gtrsim 1$.

2.4 Raman Selection Rules

A final point to highlight in this section is regarding the selection rules that the Raman coupling imposes on the signal and control fields in order to couple between the initial state and the storage state.

In order to see how these arise, from the form of C we see that

$$C \sim (\boldsymbol{d}_{12}^*.\boldsymbol{v}_s)\,(\boldsymbol{d}_{23}.\boldsymbol{v}_c), \tag{2.69}$$

where we must remove the coarse-graining of the three level model. Recalling

$$|1\rangle \to |(J=1/2,\, I)\, F=4, m_F\rangle \tag{2.70}$$
$$|2\rangle \to \left|(J=3/2,\, I)\, F', m_F'\right\rangle \tag{2.71}$$
$$|3\rangle \to \left|(J=1/2,\, I)\, F''=3, m_F''\right\rangle, \tag{2.72}$$

where $F' = 2, 3, 4, 5$, and the nuclear spin of caesium is $I = 7/2$. More explicitly then, the terms we must consider are of the form

$$\begin{aligned}\boldsymbol{d}_{12}^*.\boldsymbol{v}_s &= \langle F', m_F'| - e\boldsymbol{r}\,|F=4, m_F\rangle\,.\boldsymbol{v}_s \\ &= -e\sum_{q=-1}^{1}(-1)^q\, v_{s_{-q}}\,\langle F', m_F'|\, r_q\,|F=4, m_F\rangle\,,\end{aligned} \tag{2.73}$$

and likewise for the control, where we have expanded the vectors (rank 1 tensors) into their spherical tensor components, $q = 0, \pm 1$. Then, recalling that the quantum numbers F are a result of coupling three angular momentum components[15] and the dipole transition only addresses the *orbital* component.

Using the Wigner-Eckart theorem and decoupling the angular momentum components, one can write [143]

$$\begin{aligned}\langle F, m_F|r_q|F', m_F'\rangle &= \langle F||\boldsymbol{r}||F'\rangle\,(-1)^{F'-1+m_F}\sqrt{2F+1}\begin{pmatrix}F' & 1 & F\\ m_F' & q & -m_F\end{pmatrix} \\ \text{where}\qquad \langle F||\boldsymbol{r}||F'\rangle &= \langle J||\boldsymbol{r}||J'\rangle\,(-1)^{F'+J+1+I}\sqrt{2F')(2J+1)}\begin{Bmatrix}J & J' & 1\\ F' & F & I\end{Bmatrix},\end{aligned} \tag{2.74}$$

where the reduced matrix elements (double bars) are independent of the projection direction and the

[15] The orbital angular momentum L and electron spin S into J, and this couples with the nuclear spin I to give F.

Wigner $3j$ symbol (round braces) contains the orientation information; and the Wigner $6j$ symbol (curly braces) arises from coupling relationship of three angular momentum components (J and I into F). The value of $\langle J||\boldsymbol{r}||J'\rangle$ is constant for all the transitions and can be evaluated directly from measurement of the transition lifetime [144] and is tabulated in [143].

So then, including both steps of the transition and summing over all possible intermediate states $|F', m'_F\rangle$, and noting that we are far detuned with respect to the hyperfine splitting of the excited state, such that equal weighting can be given to the different pathways, it can be shown e.g. [145], or by evaluating the Clebsch-Gordon coefficients in Mathematica:

$$c(\boldsymbol{v}_\mathrm{s}, \boldsymbol{v}_\mathrm{c}) \propto \sum_{q,q'} \sum_{m_i,m_j} \sum_{F',m'_F} (-1)^{q+q'}\, v_{s_{-q}}\, v_{c_{-q'}}\, \langle F', m'_F|\, r_q\, |4, m_i\rangle\, \langle F', m'_F|\, r_{q'}\, |3, m_j\rangle\,, \tag{2.75}$$

using the expressions above in 2.74. In so doing one finds that the coupling is only nonzero if the control and signal field polarizations have either (i) orthogonal linear polarizations; or (ii) the same circular polarization.

Since experimentally, we will need to isolate the extracted signal from the control, we opt to use orthogonal linear polarizations, and use polarization filtering.

2.5 Summary

In this chapter:

§2.2 I derive the 1D linearised Maxwell-Bloch equations of motion for the Raman memory in the adiabatic limit. This derivation includes terms to capture the origin of the FWM process (obtaining the final result in equation 2.47)

§2.3 I discuss the *ideal* operation of the Raman memory in the absence of FWM, in which the broadband beam splitter analogy is seen. I use this picture describe the performance of a memory and highlight the parameters which determine this.

PART I

SUPPRESSION OF FOUR–WAVE MIXING IN RAMAN MEMORIES

CHAPTER 3

A CAVITY ENHANCED RAMAN MEMORY

In this chapter I discuss the first efforts of our group to address the "four–wave mixing" problem in the Raman memory protocol. The experimental project was led by D.J.S., and the theory by J.N.; I worked alongside T.F.M.C.,[1] C.Q. and P.M.L. in the lab, who all contributed to building and maintaining the optical set up.

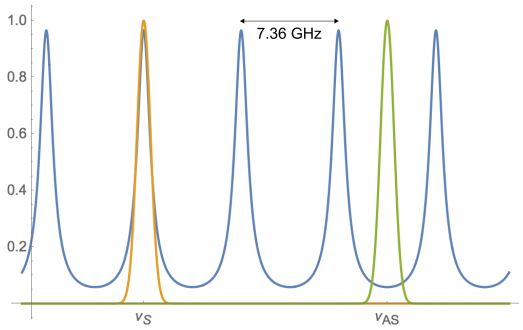
The project was already in progress when I joined the group. As well as being one of the team working on the day-to-day operation of the experiment, I designed and implemented the mechanical and thermal control systems; additionally I independently worked to simulate the atom–plus–cavity response for the characterisation of the experiment, and as a tool for future experiments. Leading on from this modelling, I assisted J.N. in developing our theoretical description of the cavity Raman memory summarised in [147], for which I had a significant role in (1) modifying the analysis to extend to the case with the large dispersive effects of introducing the atoms in the cavity outlined in the following section, and (2) exploring the consequences of this theory in practical application, which is investigated further in the next chapter.

3.1 The Idea

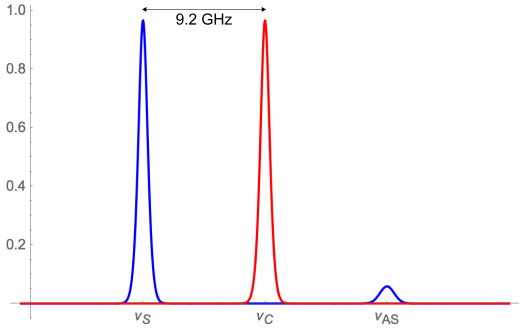
The proposal here is conceptually a simple one. Based upon our equations of motion for the Raman memory in the previous chapter, we see that in addition to the target beam splitter interaction between the signal field S and the atomic spin–wave excitation B, the control field may couple to the initially populated state, and scatter an anti–Stokes photon A – this can be interpreted as a two–mode squeezing interaction between A and B. We therefore see that a process of four–wave mixing, absorbing two control field photons, can create a pair of S and A photons, giving rise to spurious excitations in the signal mode: that is, we may retrieve photons in the same spectral, temporal and spatial modes as the signal *even when no photons were stored*.

To eliminate these spurious excitations, one could consider a number of approaches – (1) interaction selection rules; (2) phase matching; (3) Stokes shift; or (4) manipulation of density of states (DOS). Whilst selection rules for Raman coupling could be effective at suppressing noise in solid state low symmetry systems e.g. [148], it can be shown [133, 145] that this is ineffective in alkali vapours. Similarly engineering sufficient dispersion in the medium to suppress the FWM by inhibiting phase matching for the process is inefficient in vapour due to the small susceptibilities far from resonance.

[1] The final experimental results in this chapter appear in the final chapters of T.F.M.C.'s thesis [146].



(a) Low finesse cavity response matched to resonance/anti–resonance condition with signal/anti–Stokes fields.



(b) Resulting intracavity density of states for signal polarisation (blue), and control (red).

Figure 3.1: (a) Simple schematic of resonant/antiresonant cavity response superimposed with input fields (b) Resulting intracavity fields for the signal and anti-Stokes field (blue) and control field assuming that the cavity response for the orthogonal polarisation may be shifted into resonance with the control .

The difference in detuning between Stokes and anti-Stokes scattering is given by the Stokes shift, or ground state splitting $\Delta_0 = 9.2\,\mathrm{GHz}$. Since the coupling $C \propto 1/\Delta$, FWM is naturally suppressed in media with large Stokes shift. One could engineer larger Stokes shift in solid state e.g. with strain, or in vapour move closer to resonance, at the expense of achievable bandwidth, and at the risk of increased fluorescence.

Our chosen approach is to implement the memory protocol inside a cavity, suppressing the density of states of the anti-Stokes mode of the FWM, see figure 3.1, thereby reducing the coupling for the FWM process. Also, by introducing a cavity around the memory, we have the additional features that [138], firstly: the effective length of the memory interaction is increased since the signal makes multiple passes through the memory medium; and, secondly we have the cavity enhancement of the control field; together these mean that the memory could be implemented with a smaller footprint in terms of both size and power requirements than in previous bulk experiments. These features are captured by the "cooperativity" parameter of the system, which is proportional to the product of the single pass optical depth of the memory medium and the cavity finesse. The cooperativity parameter therefore encapsulates the additional cavity enhancements owing to multiple passes.

We note that, although the present, proof of principle, experiments are conducted in caesium vapour – building upon previous work, the cavity approach to suppression of FWM in a Raman memory is applicable for any medium.

3.2 Theory of Noise Suppression

Starting from the set of equations 2.47 of the form

$$(c\partial_z + \partial_t)\, S = \mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_\mathrm{s}}B - \kappa_\mathrm{s}S \tag{3.1}$$

$$(c\partial_z + \partial_t)\, A = \mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_\mathrm{a}}B^\dagger - \kappa_\mathrm{a}A \tag{3.2}$$

$$\partial_t B = -\mathrm{i}\sqrt{\frac{d\gamma}{\mathrm{L}}}\left[\frac{\Omega^*}{\Gamma_\mathrm{s}}S - \frac{\Omega}{\Gamma_\mathrm{a}}A^\dagger\right] - \left[\frac{1}{\Gamma_\mathrm{s}} + \frac{1}{\Gamma_\mathrm{a}^*}\right]|\Omega|^2 B, \tag{3.3}$$

where we have assumed for compactness that the Rabi frequency and weak-field couplings to either arm of the Λ system are approximately equal and the distinction between the two paths may be neglected – and we drop the bracketed superscripts in equation 2.47.

As before, Γ_f denotes the complex detuning for the respective fields f at the signal ($f = \mathrm{s}$) [anti–Stokes ($f = \mathrm{a}$)] frequencies. The spatial parameter z now represents the position along the optical path inside the cavity. The absorption/ dispersion for A and S due to the atomic ensemble *far from resonance*[2] is given by

$$\kappa_\mathrm{s} = \frac{d\gamma}{\tau\Gamma_\mathrm{s}}; \tag{3.4}$$

$$\kappa_\mathrm{a} = \frac{d\gamma}{\tau\Gamma_\mathrm{a}^+}, \tag{3.5}$$

where $\tau = L/c$ is the cavity roundtrip time.[3] Here $\Gamma_\mathrm{a}^+ = \gamma - \mathrm{i}(\Delta_\mathrm{a} + \Delta_0)$, since κ_f arises from coupling to the populated level. Here we neglect the decoherence of the spin wave.[4]

By inspection, we may break down these equations of motion into its functional components

- beam splitter between signal mode S and spin wave B;
- two mode squeezer between A and B;
- A.C. Stark shift $\propto |\Omega|^2$;
- complex decay term (incl. dispersion) for the propagation of S and A.

We will refer to these four ingredients throughout this thesis.

This set of equations may be solved analytically [149]; however, here we exploit the additional physics of introducing a cavity around the atomic ensemble to simplify the analysis.

In the limit of weak interaction with the atoms on a single pass of the cavity, we may relate the fields after traversing $z = L$ within the cavity to the input at $z = 0$ by Taylor expansion – see figure 3.2. Explicitly including the roundtrip dispersion/absorption and writing $\tilde{S} = \mathrm{e}^{\kappa_\mathrm{s} z/c} S$, and likewise for A, we may write:

$$\tilde{S}_{z=L} \approx \tilde{S}_{z=0} + L\partial_z \, \tilde{S}\Big|_{z=0}, \tag{3.6}$$

where using the terms from equation 3.1 the spatial derivative evaluates to

$$\begin{aligned}
L\partial_z \, \tilde{S}\Big|_{z=0} &= L\partial_z \left(\mathrm{e}^{\kappa_\mathrm{s} z/c} S\right)_{z=0} \\
&= \frac{L}{c} \, \mathrm{e}^{\kappa_\mathrm{s} z/c}\Big|_{z=0} \left(\kappa_\mathrm{s} + c\partial_z\right) S|_{z=0} \\
&= \tau \left[\kappa_\mathrm{s} S_0 + \left(\mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_\mathrm{s}} B - \kappa_\mathrm{s} S - \partial_t S\right)_{z=0}\right] \\
&= \tau \left[\mathrm{i}c\sqrt{\frac{d\gamma}{L}}\frac{\Omega}{\Gamma_\mathrm{s}} B_0 - \partial_t S_0\right] \\
&= \mathrm{i}\frac{\sqrt{d\gamma}\Omega}{\Gamma_\mathrm{s}}\sqrt{L} B_0 - \tau\partial_t S_0,
\end{aligned} \tag{3.7}$$

[2] Here, the "far detuned" condition means that inhomogeneous (Gaussian) broadening may be neglected in comparison to the homogeneous (Lorentzian) linewidth.

[3] Distinct from the "retarded time" coordinate in the previous chapter.

[4] Typical decoherence mechanisms and their timescales for these experiments are discussed further in section 4.1.2.

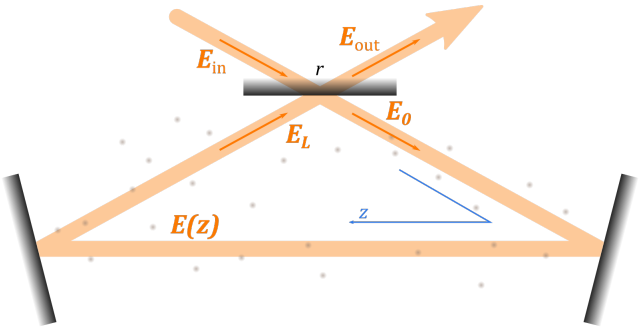


Figure 3.2: Schematic of the field E traversing around the ring cavity, with input/output coupler of reflectivity r. Here E labels the weak fields S or A.

for the signal, and analogously for the anti–Stokes. Putting these into equations 3.6 we have

$$\tilde{S}_L \approx \left[S_0 + \mathrm{i}\frac{\sqrt{d\gamma}\Omega}{\Gamma_\mathrm{s}}\sqrt{L}B_0 - \tau\partial_t S_0\right]. \tag{3.8}$$

We may then impose the cavity boundary conditions at the input coupler

$$S_0 = r\mathrm{e}^{\mathrm{i}k_\mathrm{s}L}S_L + t_r S_\mathrm{in}, \tag{3.9}$$

where $t_r = \sqrt{1-r^2}$. And $k_{f\in\{\mathrm{s,a}\}}$ are the carrier wavevectors for the respective fields. Inserting expressions for $S_L = \mathrm{e}^{-\kappa_\mathrm{s}\tau}\tilde{S}_L$ ($A_L = \mathrm{e}^{-\kappa_\mathrm{a}\tau}\tilde{A}_L$) from 3.8 (which now includes the information about the atomic susceptibility) into 3.9 then yields

$$\begin{aligned} S_0 &= r\mathrm{e}^{\mathrm{i}k_\mathrm{s}L}\mathrm{e}^{-\kappa_\mathrm{s}\tau}\tilde{S}_L + t_r S_\mathrm{in} \\ &= \zeta_\mathrm{s}\left[S_0 + \mathrm{i}\frac{\sqrt{d\gamma}\Omega}{\Gamma_\mathrm{s}}\sqrt{L}B_0 - \tau\partial_t S_0\right] + t_r S_\mathrm{in} \end{aligned} \tag{3.10}$$

in which we have grouped together the cavity dispersion and loss for the respective fields, and defined:

$$\zeta_f = \mu_f\mathrm{e}^{\mathrm{i}\phi_f} = r\varsigma\mathrm{e}^{\mathrm{i}k_f L - \kappa_f\tau}, \tag{3.11}$$

where we identify $\phi_f = k_f - \mathfrak{Im}\{\kappa_f\}$ as the total roundtrip phase accrued by the respective fields *including dispersion induced by the atomic ensemble*, and correspondingly $\mu_f = r\varsigma\mathrm{e}^{-\mathfrak{Re}\{\kappa_f\}}$ encoding the losses: from the in–coupler, atomic ensemble, and where an extra factor ς is included here by hand to account for any additional intracavity losses experienced by the fields.

Proceeding with this, and with a little rearranging, we find

$$\partial_t\left(\sqrt{\tau}S_0\right) = \frac{\zeta_\mathrm{s}-1}{\tau\zeta_\mathrm{s}}\left(\sqrt{\tau}S_0\right) + \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_\mathrm{s}}\left(\sqrt{L}B_0\right) + \frac{t_r}{\zeta_\mathrm{s}\sqrt{\tau}}S_\mathrm{in}, \tag{3.12}$$

and, by inspection:

$$\partial_t\left(\sqrt{\tau}A_0\right) = \frac{\zeta_\mathrm{a}-1}{\tau\zeta_\mathrm{a}}\left(\sqrt{\tau}A_0\right) + \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_\mathrm{a}}\left(\sqrt{L}B_0^\dagger\right) + \frac{t_r}{\zeta_\mathrm{a}\sqrt{\tau}}A_\mathrm{in}. \tag{3.13}$$

We then transform to the intracavity field amplitudes, as in [150]

$$\left.\begin{aligned} s &= \sqrt{\tau} S_0 \\ a &= \sqrt{\tau} A_0 \\ b &= \sqrt{L} B_0 \end{aligned}\right\} \tag{3.14}$$

and defining the cavity decay rates for the respective fields:

$$\frac{1}{\gamma_f} = \frac{\tau \zeta_f}{1-\zeta_f}, \tag{3.15}$$

we arrive at

$$\partial_t s = -\gamma_\mathrm{s} s + \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_\mathrm{s}} b + \frac{t_r}{\zeta_\mathrm{s}\sqrt{\tau}} S_\mathrm{in} \tag{3.16}$$

$$\partial_t a = -\gamma_\mathrm{a} a + \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_\mathrm{a}} b^\dagger + \frac{t_r}{\zeta_\mathrm{a}\sqrt{\tau}} A_\mathrm{in} \tag{3.17}$$

$$\partial_t b = -\mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\left[\frac{\Omega^*}{\Gamma_\mathrm{s}} s - \frac{\Omega}{\Gamma_\mathrm{a}} a^\dagger\right] - \left[\frac{1}{\Gamma_\mathrm{s}} + \frac{1}{\Gamma_\mathrm{a}^*}\right]|\Omega|^2 b. \tag{3.18}$$

3.2.1 Bad Cavity Limit

In the bad cavity limit, the fields incident on the cavity are *much more narrow band than the cavity linewidth*; that is, we require that $|\partial_t s| \ll |\gamma_\mathrm{s} s|$, with comparable bandwidth for the control field. In this limit, we set $\partial_t s = \partial_t a \approx 0$ in equations 3.16–3.17 and find a and b to obtain:

$$s = \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_\mathrm{s}\gamma_\mathrm{s}} b + \frac{t_r}{\gamma_\mathrm{s}\zeta_\mathrm{s}\sqrt{\tau}} S_\mathrm{in}, \tag{3.19}$$

$$a = \mathrm{i}\sqrt{\frac{d\gamma}{\tau}}\frac{\Omega}{\Gamma_\mathrm{a}\gamma_\mathrm{a}} b^\dagger + \frac{t_r}{\gamma_\mathrm{a}\zeta_\mathrm{a}\sqrt{\tau}} A_\mathrm{in}, \tag{3.20}$$

and substituting these expressions into 3.18

$$\partial_t b = \left[\frac{d\gamma}{\tau}\left(\frac{1}{\Gamma_\mathrm{s}^2\gamma_\mathrm{s}} + \frac{1}{|\Gamma_\mathrm{a}|^2\gamma_\mathrm{a}^*}\right) - \left(\frac{1}{\Gamma_\mathrm{s}} + \frac{1}{\Gamma_\mathrm{a}^*}\right)\right]|\Omega|^2 b + \mathrm{i}\frac{t_r\sqrt{d\gamma}}{\tau}\left[\frac{\Omega^*}{\zeta_\mathrm{s}\Gamma_\mathrm{s}\gamma_\mathrm{s}} S_\mathrm{in} + \frac{\Omega}{\zeta_\mathrm{a}^*\Gamma_\mathrm{a}\gamma_\mathrm{a}^*} A_\mathrm{in}^\dagger\right]. \tag{3.21}$$

We now make the coordinate transformation, in which the time parameter is defined by the Rabi frequency, thereby removing dependence on the temporal shape of the control field:

$$t \to \epsilon(t) = \frac{1}{W}\int_{-\infty}^{t} \mathrm{d}t' \left|\Omega(t')\right|^2 \tag{3.22}$$

$$\text{where } W = \int_{-\infty}^{\infty} \mathrm{d}t' \left|\Omega(t')\right|^2; \tag{3.23}$$

$$\Rightarrow \quad \partial_t = \frac{|\Omega(t)|^2}{W}\partial_\epsilon, \tag{3.24}$$

and define correspondingly normalised fields:

$$\left.\begin{aligned} \sigma &= s\sqrt{W}/\Omega = S_0\sqrt{W\tau}/\Omega \\ \alpha &= a\sqrt{W}/\Omega = A_0\sqrt{W\tau}/\Omega \end{aligned}\right\} \tag{3.25}$$

with which, the set of equations 3.19–3.21 reduce to

$$\sigma = c_{\mathrm{s}} b + p_{\mathrm{s}} \sigma_{\mathrm{in}} \tag{3.26}$$

$$\alpha = c_{\mathrm{a}} b^{\dagger} + p_{\mathrm{a}} \alpha_{\mathrm{in}} \tag{3.27}$$

$$\partial_{\epsilon} b = h b + g_{\mathrm{s}} \sigma_{\mathrm{in}} - g_{\mathrm{a}}^{*} \alpha_{\mathrm{in}}^{\dagger}, \tag{3.28}$$

where we have defined the coupling constants, c_f, and effective transmisions, p_f, for the beam–splitter [two–mode–squeezer] at the signal [anti–Stokes] frequencies

$$c_f = \frac{\mathrm{i}\bar{C}_f}{\gamma_f \sqrt{\tau}} \tag{3.29}$$

$$p_f = \frac{t_r}{\zeta_f \gamma_f \sqrt{\tau}}, \tag{3.30}$$

recalling $\bar{C}_f = \sqrt{d\gamma W}/\Gamma_f$ as the free–space Raman coupling from the previous chapter. The constant, g_f, is given by:

$$g_{\mathrm{s}} = \mathrm{i}\left(\frac{t_r}{\zeta_{\mathrm{s}}} \frac{\bar{C}_{\mathrm{s}}}{\gamma_{\mathrm{s}} \tau}\right); \qquad g_{\mathrm{a}} = \mathrm{i}\left(\frac{t_r}{\zeta_{\mathrm{a}}} \frac{\bar{C}_{\mathrm{a}}^{*}}{\gamma_{\mathrm{a}} \tau}\right), \tag{3.31}$$

noting the difference in conjugation in the appearance of $\bar{C}_{\mathrm{s}}$ and $\bar{C}_{\mathrm{a}}$. Lastly, h, which encapsulates the coupling of the memory interaction:

$$h = W\left[\frac{d\gamma}{\tau}\left(\frac{1}{\Gamma_{\mathrm{s}}^2 \gamma_{\mathrm{s}}} + \frac{1}{|\Gamma_{\mathrm{a}}|^2 \gamma_{\mathrm{a}}^{*}}\right) - \left(\frac{1}{\Gamma_{\mathrm{s}}} + \frac{1}{\Gamma_{\mathrm{a}}^{*}}\right)\right], \tag{3.32}$$

which we can re–write as

$$h = \frac{1}{\gamma_{\mathrm{s}} \tau}\left(\bar{C}_{\mathrm{s}}^2 + \frac{\gamma_{\mathrm{s}} \Gamma_{\mathrm{a}}}{\gamma_{\mathrm{a}}^{*} \Gamma_{\mathrm{a}}^{*}} \bar{C}_{\mathrm{a}}^2\right) - W\left(\frac{1}{\Gamma_{\mathrm{s}}} + \frac{1}{\Gamma_{\mathrm{a}}^{*}}\right) \tag{3.33}$$

$$= \left(C_{\mathrm{s}} + \frac{\gamma_{\mathrm{s}}}{\gamma_{\mathrm{a}}^{*}} \mathrm{e}^{2\mathrm{i}\arg(\Gamma_{\mathrm{a}})} C_{\mathrm{a}}\right) - W\left(\frac{1}{\Gamma_{\mathrm{s}}} + \frac{1}{\Gamma_{\mathrm{a}}^{*}}\right), \tag{3.34}$$

where the Raman coupling $\bar{C}_f$ introduced in the previous chapter is modified by the cavity to become

$$\bar{C}_f \to C_f = \frac{\sqrt{\mathcal{C}\gamma W}}{\Gamma_f} = \frac{\bar{C}_f}{\sqrt{\gamma_s \tau}}; \tag{3.35}$$

$$\text{where} \qquad \mathcal{C} = \frac{d}{\gamma_{\mathrm{s}} \tau}. \tag{3.36}$$

The parameter $\mathcal{C}$ represents the enhanced optical depth, or *cooperativity* of the cavity, *for the signal field*. Note that on resonance with the cavity, $\phi_{\mathrm{s}} = 0$, this is equivalent to interacting with a medium with an effective length enhanced proportional to the cavity finesse, $\mathcal{F}_{\mathrm{s}}$, when the control field Ω is not present: that is $L \to L \times \mathcal{F}_{\mathrm{s}}/\pi$.

We note that the fields defined by the set of equations 3.26–3.28 are related to the original fields by unitary transformation, hence the operators $o \in \{\sigma, \alpha, \beta\}$ obey the bosonic commutation relations $\left[o, o^{\dagger}\right] = 1$.

To find the output optical fields, we again impose analogous boundary conditions as equation 3.9

$$S_{\mathrm{out}} = t_r \mathrm{e}^{\mathrm{i}k_{\mathrm{s}} L} S_L - r S_{\mathrm{in}}, \tag{3.37}$$

and noting that in the limit of weak single–pass coupling $\tilde{S}_L \approx (\zeta_s/r)\tilde{S}_0$. The output signal field in a given time–bin[5] j is thus given by:

$$\sigma_{\text{out},j}(\epsilon) = \mathrm{i}\chi\bar{C}_\text{s}\frac{\zeta_\text{s}}{t_r}b_j(\epsilon) + (\chi - r)\sigma_{\text{in},j}(\epsilon), \tag{3.38}$$

where we define the cavity roundtrip transmission function at the signal frequency:

$$\chi = \frac{t_r^2}{r}\frac{\zeta_\text{s}}{1-\zeta_\text{s}} \tag{3.39}$$

$$= \frac{t_r^2}{r} \times \frac{1}{\gamma_s\tau}. \tag{3.40}$$

Integrating equation 3.28, we see that $b_j(\epsilon)$ may be expressed as:

$$b_j(\epsilon) = b_j(0)\mathrm{e}^{h\epsilon} + \int_0^1 \mathrm{d}\epsilon' M_c(\epsilon,\epsilon')\left[g_\text{s}\sigma_{\text{in},j}(\epsilon') + g_\text{a}^*\alpha^\dagger_{\text{in},j}(\epsilon')\right] \tag{3.41}$$

$$\text{where} \qquad M_c(\epsilon,\epsilon') = \Theta(\epsilon,\epsilon')M(\epsilon,\epsilon') \tag{3.42}$$

$$\text{and} \qquad M(\epsilon,\epsilon') = \exp\left[h\left(\epsilon-\epsilon'\right)\right]. \tag{3.43}$$

Here, Θ is the Heaviside step–function which imposes the temporal causality of the cavity response function M.

Inserting 3.41 into 3.38, at time ϵ

$$\sigma_{\text{out},j}(\epsilon) = K \times b_j(\epsilon) + \left[\chi - r\right]\sigma_{\text{in},j}(\epsilon); \qquad \left(\text{with} \quad K = \mathrm{i}\chi\bar{C}_\text{s}\frac{\zeta_\text{s}}{t_r}\right) \tag{3.44}$$

$$= Kb_j(0)\mathrm{e}^{h\epsilon} + \int_0^1 \mathrm{d}\epsilon' K M_c(\epsilon,\epsilon')\left\{g_\text{s}\sigma_{\text{in},j}(\epsilon') - g_\text{a}^*\alpha^\dagger_{\text{in},j}(\epsilon')\right\} + (\chi - r)\sigma_{\text{in},j}(\epsilon)$$

$$= Kb_j(0)\mathrm{e}^{h\epsilon} + \int_0^1 \mathrm{d}\epsilon' \left[\left\{KM_c(\epsilon,\epsilon')g_\text{s} + (\chi-r)\delta(\epsilon-\epsilon')\right\}\sigma_{\text{in},j}(\epsilon') - KM_c(\epsilon,\epsilon')g_\text{a}^*\alpha^\dagger_{\text{in},j}(\epsilon')\right]. \tag{3.45}$$

3.2.1(a) Storage

In the first time–bin $j = 0$ when the signal is sent into the memory; if the system is initialised with no spin wave, $b_0(0) = 0$, these reduce to

$$b_0(\epsilon) = \int_0^1 \mathrm{d}\epsilon' M_c(\epsilon,\epsilon')\left[g_\text{s}\sigma_{\text{in},0}(\epsilon') + g_\text{a}^*\alpha^\dagger_{\text{in},0}(\epsilon')\right] \tag{3.46}$$

$$\sigma_{\text{out},0}(\epsilon) = \int_0^1 \mathrm{d}\epsilon' \left[M^{[0]}_\text{Ram}(\epsilon,\epsilon')\sigma_{\text{in},0}(\epsilon') + M^{[0]}_\text{FWM}(\epsilon,\epsilon')\alpha^\dagger_{\text{in},0}(\epsilon')\right]. \tag{3.47}$$

The square bracketed superscript $[n]$ on the beam–splitter (Ram) and four–wave–mixing (FWM) transfer functions indicate the contribution from the n^{th} preceding time bin. Explicitly evaluating the integral kernels using equation 3.45, we see that we can write:

$$M^{[0]}_\text{Ram}(\epsilon,\epsilon') = -\chi C_\text{s}^2 M_c(\epsilon,\epsilon') + (\chi - r)\delta(\epsilon-\epsilon') \tag{3.48}$$

$$M^{[0]}_\text{FWM}(\epsilon,\epsilon') = -\chi C_\text{s}C_\text{a}xM_c(\epsilon,\epsilon'). \tag{3.49}$$

[5] Where each "bin" index labels successive control pulses incident on the memory, and ϵ describes the normalised time elapsed *in that bin*.

Starting to think about noise suppression we note that the anti–Stokes coupling C_{a} always appears with an additional factor x with respect to the signal coupling, where

$$x = \frac{\zeta_{\mathrm{s}}\gamma_{\mathrm{s}}}{(\zeta_{\mathrm{a}}\gamma_{\mathrm{a}})^*} = \frac{1-\zeta_{\mathrm{s}}}{1-\zeta_{\mathrm{a}}^*}, \tag{3.50}$$

and we thus term this the *noise suppression factor*.

3.2.1(b) Retrieval

Similarly, in the subsequent time–bin, the retrieval interaction, the input spin wave is taken as the excitation created in the read–in bin,[6] $b_1(0) = b_0(1)$, and there is no incident signal, $\sigma_{\mathrm{in},1} = 0$.

For now, assuming an identical control pulse for both read–in and read–out interactions, the resulting spin–wave is

$$\begin{aligned} b_1(\epsilon) &= \mathrm{e}^{h\epsilon} \times \left\{ \int_0^1 \mathrm{d}\epsilon' M_c(1,\epsilon') \left[g_{\mathrm{s}}\sigma_{\mathrm{in},0}(\epsilon') + g_{\mathrm{a}}^*\alpha_{\mathrm{in},0}^\dagger(\epsilon') \right] \right\} + \int_0^1 \mathrm{d}\epsilon' M_c(\epsilon,\epsilon') \left[g_{\mathrm{s}}\sigma_{\mathrm{in},1}(\epsilon') + g_{\mathrm{a}}^*\alpha_{\mathrm{in},1}^\dagger(\epsilon') \right] \\ &= \int_0^1 \mathrm{d}\epsilon'\, \mathrm{e}^{h\epsilon} \times 1 \times \mathrm{e}^{h(1-\epsilon')} \left[g_{\mathrm{s}}\sigma_{\mathrm{in},0}(\epsilon') + g_{\mathrm{a}}^*\alpha_{\mathrm{in},0}^\dagger(\epsilon') \right] + \int_0^1 \mathrm{d}\epsilon' \left[0 + M_c(\epsilon,\epsilon') g_{\mathrm{a}}^*\alpha_{\mathrm{in},1}^\dagger(\epsilon') \right] \\ &= \int_0^1 \mathrm{d}\epsilon' \left\{ \mathrm{e}^h M(\epsilon,\epsilon') \left[g_{\mathrm{s}}\sigma_{\mathrm{in},0}(\epsilon') + g_{\mathrm{a}}^*\alpha_{\mathrm{in},0}^\dagger(\epsilon') \right] + \left[M_c(\epsilon,\epsilon') g_{\mathrm{a}}^*\alpha_{\mathrm{in},1}^\dagger(\epsilon') \right] \right\}. \end{aligned} \tag{3.51}$$

Hence, we can write the signal field as:

$$\sigma_{\mathrm{out},1} = \int_0^1 \mathrm{d}\epsilon' \left[M_{\mathrm{Ram}}^{[1]}(\epsilon,\epsilon')\sigma_{\mathrm{in},0}(\epsilon') + M_{\mathrm{FWM}}^{[1]}(\epsilon,\epsilon')\alpha_{\mathrm{in},0}^\dagger(\epsilon') + M_{\mathrm{FWM}}^{[0]}(\epsilon,\epsilon')\alpha_{\mathrm{in},1}^\dagger(\epsilon') \right], \tag{3.52}$$

where again, evaluation of the new kernels yields:

$$M_{\mathrm{Ram}}^{[1]} = -\chi C_{\mathrm{s}}^2 \mathrm{e}^h M; \tag{3.53}$$

$$M_{\mathrm{FWM}}^{[1]} = -\chi C_{\mathrm{s}} C_{\mathrm{a}} x \mathrm{e}^h M. \tag{3.54}$$

We highlight here that, from the form of 3.52, the mapping from the input in the signal mode to the retrieved signal is desribed by the Green's function $M_{\mathrm{Ram}}^{[1]} \propto M$. This function is seperable in ϵ and ϵ', where if one looks more closely:

$$\int_0^1 \mathrm{d}\epsilon' \left[M_{\mathrm{Ram}}^{[1]}(\epsilon,\epsilon')\sigma_{\mathrm{in},0}(\epsilon') \right] = -\chi C_{\mathrm{s}}^2 \mathrm{e}^{h(1+\epsilon)} \int_0^1 \mathrm{d}\epsilon' \left[\mathrm{e}^{-h\epsilon'}\sigma_{\mathrm{in},0}(\epsilon') \right], \tag{3.55}$$

in which we may interpret that $\mathrm{e}^{-h\epsilon(t)}$ determines the temporal–mode which is addressed by the memory: *any mode orthogonal to this will not be coupled.* This is a key feature of the cavity memory, which we will return to later in this thesis.

[6] Neglecting decoherence during the storage time - which is assumed to occur on much longer timescales.

3.2.1(c) Subsequent Retrieval

By extension to equation 3.52, if we have a train of identical control pulses, with the memory perfectly initialised $b_0(0) = 0$ and $\sigma_{\text{in},j} = 0$ for $j > 0$; we can always write the output signal field in the form:

$$\sigma_{\text{out},j}(\epsilon) = \int_0^1 \mathrm{d}\epsilon' \left[M^{[j]}_{\text{Ram}}(\epsilon, \epsilon')\sigma_{\text{in},0}(\epsilon') + \sum_{j'=0}^{j} \left(M^{[j']}_{\text{FWM}}(\epsilon, \epsilon')\alpha^\dagger_{\text{in},j-j'}(\epsilon') \right) \right]. \tag{3.56}$$

3.2.2 Performance

With this framework, we may begin to model the behaviour of the Raman memory in a cavity. The principal metrics we choose to consider here are the memory efficiency; and the statistics of the retrieved state in order to evaluate how well the cavity Raman memory will preserve the quantum character of the input state.

3.2.2(a) Efficiency & SNR

To evaluate the memory efficiency, we consider the expectation values of excitation number of the signal field ($f = \text{s}$) incident upon–/ retrieved from– ($i \in \{\text{in}, \text{out}\}$) the memory in a given time–bin (j); that is

$$N_{i,j} = \left\langle \int_{-\infty}^{\infty} \mathrm{d}t' \; S^\dagger_{i,j}(t') S_{i,j}(t') \right\rangle = \left\langle \int_0^1 \mathrm{d}\epsilon' \; \sigma^\dagger_{i,j}(\epsilon')\sigma_{i,j}(\epsilon') \right\rangle \qquad \text{(by unitarity)} \tag{3.57}$$

where the expectation value for some operator $\mathcal{O}$ acting on the system, $\langle \mathcal{O} \rangle$ indicates $\text{Tr}\left[\rho \mathcal{O}\right]$ where ρ is the density matrix defining the collective system. Expanding this out for the output signal mode:

$$\begin{aligned} \sigma^\dagger_{\text{out},j}(\epsilon)\sigma_{\text{out},j}(\epsilon) = \int_0^1 \mathrm{d}\epsilon' \int_0^1 \mathrm{d}\epsilon'' \Bigg[& \sigma^\dagger_{\text{in},0}(\epsilon'') M^{[j]\,*}_{\text{Ram}}(\epsilon, \epsilon'') M^{[j]}_{\text{Ram}}(\epsilon, \epsilon')\sigma_{\text{in},0}(\epsilon') \\ & + \sum_{j'=0}^{j} \left(\alpha_{\text{in},j-j'}(\epsilon'') M^{[j']\,*}_{\text{FWM}}(\epsilon, \epsilon'') M^{[j']}_{\text{FWM}}(\epsilon, \epsilon')\alpha^\dagger_{\text{in},j-j'}(\epsilon') \right) \\ & + \dots \Bigg], \end{aligned} \tag{3.58}$$

where we have only retained normally/anti–normally ordered terms in the *same mode* and the *same time–bin*, which can give rise to non–zero contributions when the expectation values are evaluated, and have discarded any cross–terms. We see that, whilst the input of the anti–Stokes field is in the vacuum state, equation 3.57 will give rise to terms of the form $\alpha\alpha^\dagger$, thus giving rise to noise in the interaction; e.g. for the usual ladder operators a, we can write $\langle aa^\dagger \rangle = n + 1$ for a Fock state $|n\rangle$.

For subsequent analysis, we adopt an abbreviated vector/matrix notation, in which we represent the fields as vectors and the two–dimensional functions as matrices:[7]

$$v(y) \to \vec{v} = \begin{pmatrix} v(y_0) \\ v(y_1) \\ \vdots \\ v(y_{N-1}) \end{pmatrix}; \quad K(x, y) \to \boldsymbol{M} = \begin{pmatrix} K(x_0, y_0) & K(x_0, y_1) & \cdots & K(x_0, y_{N-1}) \\ K(x_1, y_0) & K(x_1, y_1) & \cdots & K(x_1, y_{N-1}) \\ \vdots & \vdots & \ddots & \vdots \\ K(x_{N-1}, y_0) & K(x_{N-1}, y_1) & \cdots & K(x_{N-1}, y_{N-1}) \end{pmatrix},$$

[7] For later numerical evaluation, the choice of discretisation for the spatial/temporal coordinates will become important; however, for the immediate analysis, we still consider these as continuous functions.

and the typical integrals that will arise can be written

$$\int \mathrm{d}y\ K_1(x,y)K_2^*(z,y) = \boldsymbol{K}_1\boldsymbol{K}_2^\dagger$$
$$\int \mathrm{d}y\ K(y,y) = \mathrm{Tr}_\epsilon[\boldsymbol{K}]$$
$$\int \mathrm{d}y\ K(x,y)v(y) = \boldsymbol{K}\vec{v}$$
$$\int \mathrm{d}y\ v^*(y)v(y) = \vec{v}^\dagger\cdot\vec{v} = \mathrm{Tr}_\epsilon\left[\vec{v}\cdot\vec{v}^\dagger\right].$$

To quantify the performance, we initially consider the efficiency and signal–to–noise ratio of the memory in the first read–out time bin:

$$\eta_{\mathrm{mem}} = \frac{N_{\mathrm{out},1} - N_{\mathrm{out},1}|_{\text{no input}}}{N_{\mathrm{in},0}} \tag{3.59}$$
$$\mathrm{SNR} = \frac{N_{\mathrm{out},1} - N_{\mathrm{out},1}|_{\text{no input}}}{N_{\mathrm{out},1}|_{\text{no input}}}, \tag{3.60}$$

where the noise floor $N_{\mathrm{out},j}|_{\text{no input}}$ is the number of output photons in time–bin j when no signal is sent into the memory.[8]

Let us now consider sending a signal with a temporal–mode $\vec{\psi}_{\mathrm{in}}(\epsilon)$ into the memory; and seek to find the number of photons in the emerging field. From the form of 3.56, we see this acts upon the input signal/anti–Stokes modes in all preceeding bins, if we therefore write the input state as:

$$|\vec{\Psi}_{\mathrm{in}}(\epsilon)\rangle \sim \bigotimes_{j'=0}^{j}\left[|\vec{\psi}_{\mathrm{s},j'}(\epsilon)\rangle \otimes |\vec{\psi}_{\mathrm{a},j'}(\epsilon)\rangle\right] \tag{3.61}$$
$$\boldsymbol{\rho}_{\mathrm{in}}(\epsilon) = |\vec{\Psi}_{\mathrm{in}}(\epsilon)\rangle\langle\vec{\Psi}_{\mathrm{in}}(\epsilon)|, \tag{3.62}$$

where $|\psi_{i,j}(\epsilon)\rangle = \vec{\psi}_{i,j}(\epsilon)\,|\varphi_{i,j}\rangle$ is the input state of mode $i \in \{\mathrm{s},\mathrm{a}\}$ in time–bin j; it is helpful to separate out the temporal–mode function of the state $\vec{\psi}_{i,j}(\epsilon)$ (normalised to unity), and encode the rest of the state as $|\varphi_{i,j}\rangle$, satisfying $\langle\varphi_{i,j}|\varphi_{i',j'}\rangle = \delta_{i,i'}\delta_{j,j'}$. Note that the only non–vacuum state input is in the signal mode, in time bin $j = 0$, whilst for all others:

$$|\varphi_{\mathrm{s},j'}\rangle = |\mathrm{vac.}\rangle \qquad \forall j' > 0 \tag{3.63}$$
$$|\varphi_{\mathrm{a},j'}\rangle = |\mathrm{vac.}\rangle \qquad \forall j'. \tag{3.64}$$

Expanding out $N_{\mathrm{out},j}$ term–by–term, as in equation 3.58, and using the properties of the creation/annihlation operators, the expectation values we will have contributions from:

$$\left\langle \int_{-\infty}^{+\infty} \mathrm{d}\epsilon'\ \sigma^\dagger_{\mathrm{in},j}(\epsilon')\sigma_{\mathrm{in},j}(\epsilon') \right\rangle = \begin{cases} N_{\mathrm{in},0} & \text{if } j = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{3.65}$$
$$\left\langle \int_{-\infty}^{+\infty} \mathrm{d}\epsilon'\ \alpha_{\mathrm{in},j'}(\epsilon')\alpha^\dagger_{\mathrm{in},j'}(\epsilon') \right\rangle = 1, \tag{3.66}$$

[8] Here, by "memory" I mean, the complete "atoms plus control field" system.

whilst all other terms act only once on each mode, and evaluate to zero:

$$\begin{aligned}
\left\langle \int_{-\infty}^{+\infty} \mathrm{d}\epsilon'\ \sigma_{\mathrm{in},j}(\epsilon')\alpha_{\mathrm{in},j'}(\epsilon') \right\rangle &= \int_{-\infty}^{+\infty} \mathrm{d}\epsilon'\ \left\langle \psi_{\mathrm{s},j}(\epsilon') \,\middle|\, \sigma_{\mathrm{in},j}(\epsilon') \,\middle|\, \psi_{\mathrm{s},j}(\epsilon') \right\rangle \left\langle \psi_{\mathrm{a},j'} \,\middle|\, \alpha_{\mathrm{in},j'}(\epsilon') \,\middle|\, \psi_{\mathrm{a},j'}(\epsilon') \right\rangle = 0, \\
\left\langle \int_{-\infty}^{+\infty} \mathrm{d}\epsilon'\ \alpha^\dagger_{\mathrm{in},j}(\epsilon')\sigma^\dagger_{\mathrm{in},j'}(\epsilon') \right\rangle &= \quad \ldots \quad = 0, \\
\left\langle \int_{-\infty}^{+\infty} \mathrm{d}\epsilon'\ \alpha_{\mathrm{in},j}(\epsilon')\alpha^\dagger_{\mathrm{in},j'\neq j}(\epsilon') \right\rangle &= \quad \ldots \quad = 0.
\end{aligned} \tag{3.67}$$

With this in mind, when writing this out more explicitly, we find the only non–zero contributions to the photon number are

1. acting on the signal mode in the storage time–bin $j' = 0$ (as in equation 3.65):

$$\begin{aligned}
n_{\mathrm{s}}^{[j]} &= \int_0^1 \mathrm{d}\epsilon \left\langle \psi_{\mathrm{s},0}(\epsilon) \,\middle|\, \int_0^1 \mathrm{d}\epsilon'' \sigma^\dagger_{\mathrm{in},0}(\epsilon'') M^{[j]\,*}_{\mathrm{Ram}}(\epsilon,\epsilon'') \int_0^1 \mathrm{d}\epsilon' M^{[j]}_{\mathrm{Ram}}(\epsilon,\epsilon')\sigma_{\mathrm{in},0}(\epsilon') \,\middle|\, \psi_{\mathrm{s},0}(\epsilon) \right\rangle \\
&\to N_{\mathrm{in},0}\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[j]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{s},0}}\cdot\vec{\psi_{\mathrm{s},0}}^{\,\dagger}\boldsymbol{M}^{[j]\,\dagger}_{\mathrm{Ram}}\right],
\end{aligned} \tag{3.68}$$

and similarly:

2. acting on the anti–Stokes mode in all time–bins $j' \leq j$ (as in equation 3.66)

$$\begin{aligned}
n_{\mathrm{a}}^{[j']} &= \int_0^1 \mathrm{d}\epsilon \left\langle \psi_{\mathrm{a},j-j'}(\epsilon) \,\middle|\, \int_0^1 \mathrm{d}\epsilon' \int_0^1 \mathrm{d}\epsilon''\ \alpha_{\mathrm{in},j-j'}(\epsilon'') M^{[j']\,*}_{\mathrm{FWM}}(\epsilon,\epsilon'') M^{[j']}_{\mathrm{FWM}}(\epsilon,\epsilon')\alpha^\dagger_{\mathrm{in},j-j'}(\epsilon') \,\middle|\, \psi_{\mathrm{a},j-j'}(\epsilon) \right\rangle \\
&\to \mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[j']}_{\mathrm{FWM}}\boldsymbol{M}^{[j']\dagger}_{\mathrm{FWM}}\right].
\end{aligned} \tag{3.69}$$

So, all together we may write:

$$N_{\mathrm{out},j} = n_{\mathrm{s}}^{[j]} + \sum_{j'=0}^{j} n_{\mathrm{a}}^{[j']} \to \mathrm{Tr}_\epsilon[\boldsymbol{P}_j] \tag{3.70}$$

$$\text{where} \qquad \boldsymbol{P}_j = N_{\mathrm{in},0}\boldsymbol{M}^{[j]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[j]\,\dagger}_{\mathrm{Ram}} + \sum_{k=0}^{j}\boldsymbol{M}^{[k]}_{\mathrm{FWM}}\boldsymbol{M}^{[k]\,\dagger}_{\mathrm{FWM}}. \tag{3.71}$$

In the first retrieval time–bin then, using equations 3.49, 3.53, 3.54

$$N_{\mathrm{out},1} = \mathrm{Tr}_\epsilon\left[N_{\mathrm{in},0}\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\,\dagger}_{\mathrm{Ram}} + \boldsymbol{M}^{[0]}_{\mathrm{FWM}}\boldsymbol{M}^{[0]\,\dagger}_{\mathrm{FWM}} + \boldsymbol{M}^{[1]}_{\mathrm{FWM}}\boldsymbol{M}^{[1]\,\dagger}_{\mathrm{FWM}}\right] \tag{3.72}$$

$$= |\chi C_{\mathrm{s}}|^2 \left\{ N_{\mathrm{in},0}\,|C_{\mathrm{s}}|^2\,\mathrm{e}^{h+h^*}\mathrm{Tr}_\epsilon\left[\boldsymbol{M}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^\dagger\right] + |C_{\mathrm{a}}x|^2\left(\mathrm{Tr}_\epsilon\left[\boldsymbol{M_c}\boldsymbol{M_c}^\dagger\right] + \mathrm{e}^{h+h^*}\mathrm{Tr}_\epsilon\left[\boldsymbol{M}\boldsymbol{M}^\dagger\right]\right)\right\}. \tag{3.73}$$

Again, we look at the integrals, term–by–term. Let us first define the (real) dimensionless parameters:

$$\xi = -\left(h + h^*\right) = -2\,\mathfrak{Re}\left\{C_{\mathrm{s}} + \frac{\gamma_{\mathrm{s}}}{\gamma_{\mathrm{a}}^*}\mathrm{e}^{2\mathrm{i}\arg(\Gamma_{\mathrm{a}})}C_{\mathrm{a}}\right\} + 2W\,\mathfrak{Re}\left\{\frac{1}{\Gamma_{\mathrm{s}}} + \frac{1}{\Gamma_{\mathrm{a}}^*}\right\}, \tag{3.74}$$

where we recall h encodes the coupling strength in the equations of motion; and, define:

$$E = \int_0^1 \mathrm{d}\epsilon'\ \mathrm{e}^{-\xi\epsilon'} = \frac{1}{\xi}\left(1 - \mathrm{e}^{-\xi}\right). \tag{3.75}$$

Firstly let us consider the coherent part of the memory interaction – that is, the part mapping input to output in the signal mode:

$$\begin{aligned}
\mathrm{e}^{h+h^*}\mathrm{Tr}_\epsilon\Big[\boldsymbol{M}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^\dagger\Big] &= \mathrm{e}^{h+h^*}\int_0^1\mathrm{d}\epsilon\,\left[\int_0^1\mathrm{d}\epsilon''\psi_{\mathrm{in}}^*(\epsilon'')M^*(\epsilon,\epsilon'')\int_0^1\mathrm{d}\epsilon' M(\epsilon,\epsilon')\psi_{\mathrm{in}}(\epsilon')\right]\\
&= \mathrm{e}^{h+h^*}\int_0^1\mathrm{d}\epsilon\,\left[\int_0^1\mathrm{d}\epsilon'\,\mathrm{e}^{h(\epsilon-\epsilon')}\psi_{\mathrm{in}}(\epsilon')\int_0^1\mathrm{d}\epsilon''\,\mathrm{e}^{h^*(\epsilon-\epsilon'')}\psi_{\mathrm{in}}^*(\epsilon'')\right]\\
&= \mathrm{e}^{-\xi}\left(\int_0^1\mathrm{d}\epsilon\,\mathrm{e}^{-\xi\epsilon}\right)\left|\int_0^1\mathrm{d}\epsilon'\,\mathrm{e}^{-h\epsilon'}\psi(\epsilon')\right|^2\\
&= \mathrm{e}^{-\xi}E\left|\int_0^1\mathrm{d}\epsilon'\,\mathrm{e}^{-h\epsilon'}\psi(\epsilon')\right|^2.
\end{aligned}\tag{3.76}$$

Let us define μ so as to quantify the normalised overlap of the signal mode with the cavity response

$$\begin{aligned}
\mu &= \frac{\int_0^1\mathrm{d}\epsilon'\,\mathrm{e}^{-h\epsilon'}\psi(\epsilon')}{\sqrt{\int_0^1\mathrm{d}\epsilon'\;|\mathrm{e}^{-h\epsilon'}|^2}\sqrt{\int_0^1\mathrm{d}\epsilon''\;|\psi(\epsilon'')|^2}} \qquad \left(\text{recalling }\int_0^1\mathrm{d}\epsilon''\;|\psi(\epsilon'')|^2=1\right)\\
&= \left(\mathrm{e}^{\xi}E\right)^{-1/2}\int_0^1\mathrm{d}\epsilon'\,\mathrm{e}^{-h\epsilon'}\psi(\epsilon'),
\end{aligned}\tag{3.77}$$

such that, when ψ is optimally matched to the cavity response, μ is maximal at unity. With this, equation 3.76 reduces to

$$\mathrm{e}^{h+h^*}\mathrm{Tr}_\epsilon\Big[\boldsymbol{M}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^\dagger\Big] = E^2\,|\mu|^2\,.\tag{3.78}$$

Then for the retrieval of noise photons into the signal mode:

$$\begin{aligned}
\mathrm{e}^{-\xi}\mathrm{Tr}_\epsilon\Big[\boldsymbol{M}\boldsymbol{M}^\dagger\Big] &= \mathrm{e}^{-\xi}\int_0^1\mathrm{d}\epsilon'\int_0^1\mathrm{d}\epsilon''\;M(\epsilon'',\epsilon')^*M(\epsilon'',\epsilon') \qquad \left(M(\epsilon_0,\epsilon_1)=\mathrm{e}^{h(\epsilon_0-\epsilon_1)}\right)\\
&= \mathrm{e}^{-\xi}\left[\int_0^1\mathrm{d}\epsilon''\,\mathrm{e}^{-\xi\epsilon''}\right]\left[\int_0^1\mathrm{d}\epsilon'\,\mathrm{e}^{+\xi\epsilon'}\right] \qquad (\xi=-(h+h^*))\\
&= \frac{1}{\xi^2}\left(1-2\mathrm{e}^{-\xi}+\mathrm{e}^{-2\xi}\right),
\end{aligned}\tag{3.79}$$

and its causal counterpart, for interactions within the same time–bin:

$$\begin{aligned}
\mathrm{Tr}_\epsilon\Big[\boldsymbol{M_c}\boldsymbol{M_c}^\dagger\Big] &= \int_0^1\mathrm{d}\epsilon'\int_0^1\mathrm{d}\epsilon''\;\mathrm{e}^{h(\epsilon''-\epsilon')}\Theta(\epsilon''-\epsilon')\mathrm{e}^{h^*(\epsilon''-\epsilon')}\Theta(\epsilon''-\epsilon')\\
&= \int_0^1\mathrm{d}\epsilon''\,\mathrm{e}^{-\xi\epsilon''}\int_0^{\epsilon''}\mathrm{d}\epsilon'\,\mathrm{e}^{+\xi\epsilon'}\\
&= \frac{1}{\xi}\left[1+\frac{1}{\xi}\left(\mathrm{e}^{-\xi}-1\right)\right],
\end{aligned}\tag{3.80}$$

and so together, we may write

$$\mathrm{Tr}_\epsilon\Big[\boldsymbol{M_c}\boldsymbol{M_c}^\dagger\Big]+\mathrm{e}^{-\xi}\mathrm{Tr}_\epsilon\Big[\boldsymbol{M}\boldsymbol{M}^\dagger\Big] = g\left(\xi\right),\tag{3.81}$$

in which we have defined:

$$g\left(\xi\right)=\frac{1}{\xi}\left[1-\mathrm{e}^{-\xi}\left(\frac{1-\mathrm{e}^{-\xi}}{\xi}\right)\right]=\frac{1}{\xi}\left(1-\mathrm{e}^{-\xi}E\right).\tag{3.82}$$

Input state	Incident $g^{(2)}$	Output $g^{(2)}$
No input	–	$1.62 \pm .04$
Coherent state	$1.01 \pm .01$	$1.69 \pm .02$
Single photon	$0.016 \pm .004$	$1.59 \pm .03$

Table 3.1: Experimentally measured $g^{(2)}$ in the free-space Raman memory [118]. ($N_{\text{in}} \approx \eta_{\text{herald}} = .25$)

And finally, putting all of these together, into equation 3.73, we find that

$$N_{\text{out},1} = |\chi C_{\text{s}}|^2 \left\{ N_{\text{in},0} \, |C_{\text{s}}|^2 \, E^2 \, |\mu|^2 + |C_{\text{a}} x|^2 \, g(\xi) \right\}, \tag{3.83}$$

and recalling the definitions 3.59 and 3.60 one can simply write down expressions for the predicted memory efficiency and signal to noise:

$$\eta_{\text{mem}} = \left| \chi C_{\text{s}}^2 E \mu \right|^2 \tag{3.84}$$

$$\text{SNR} = N_{\text{in},0} |\mu|^2 \left| \frac{C_{\text{s}}}{C_{\text{a}} x} \right|^2 \frac{E^2}{g(\xi)}. \tag{3.85}$$

The prefactors here can be reasonably intuitively interpreted – whereby the desired interaction (the prefactor to $N_{\text{in},0}$) depends on

- E Contribution from the Green's function linking the input signal $\sigma_{\text{in},0}$ to the output $\sigma_{\text{out},1}$, via: $E^2 = \text{Tr}_\epsilon \left[\boldsymbol{M} \vec{\psi_{\text{in}}} \cdot \vec{\psi_{\text{in}}}^{\dagger} \boldsymbol{M}^\dagger \right] / |\mu|^2$
- μ Overlap of input signal with optimal memory mode
- C_{s} Cavity enhanced Raman coupling of the signal field,

whilst the spurious part:

- $g(\xi)$ Total contribution from the FWM Green's functions mapping $\alpha_{j'}$ in the storage and retrieval bins, to the final retrieved signal: $g = \text{Tr}_\epsilon \left[\boldsymbol{M_c} \boldsymbol{M_c}^\dagger \right] + \text{e}^{h+h^*} \text{Tr}_\epsilon \left[\boldsymbol{M} \boldsymbol{M}^\dagger \right]$
- $C_{\text{a}} x$ Cavity (suppressed) Raman coupling of the anti-Stokes field.

3.2.2(b) Autocorrelation

Having considered the *memory* element of our quantum memory, we next address the *quantum* part.

To verfy the quantum statistics of the retreived state compared to that which we attempt to store, we require a metric to access the quantum statistics of the input state. For the application of the storage of single photons, we consider the heralded $g^{(2)}$ autocorrelation at the input and output of the memory.

This analysis follows on from the measurement of the $g^{(2)}$ of stored single photons from a PPKTP waveguide SPDC source interfaced with a Raman memory [118] analysis of the $g^{(2)}$ in various configurations are tabulated in 3.1.

It is seen that comparing single photon level coherent light and heralded single photons, the incident state values are near to the ideal values of 1 and 0, respectively. The retrieved state $g^{(2)}$ increases in both cases towards that of thermal light ($g^{(2)} = 2$), where the reduction of the $g^{(2)}$ of the retrieved state after storing single photons with respect to coherent input or noise originates from the non-classical statistics of the input light.

The increase of the $g^{(2)}$ towards thermal statistics can intuitively be understood as the noise photons introduced into the signal mode from a two-mode squeezed state across A and S as seen from the equations of motion derived in the previous chapter. After the A part of this is discarded, the spurious photons in the signal mode are left in a thermal state,

The experimental results presented by Michelberger *et al.* in [118] are consistent with the presence of FWM noise polluting the retrieved light in the signal mode as a coherent part of the memory interaction, rather than an admixture with incoherent noise. The $g^{(2)}$ is sensitive to this and so this is a useful metric for us to consider when modelling the performance of the cavity Raman memory.

To clarify this distinction: one could imagine some independent noise process producing photons indistinguishable with the signal, and combining these on a beam-splitter [151]. The output state of this would yield the combined photon statistics given by a superposition of FWM noise and that of the retrieved signal. This is in contrast to the noise photons being introduced by a coherent part of the same light-matter interaction process.

Moving forward to model this in the cavity memory, we may write down the expression for the $g^{(2)}$ of the output mode as the integral over time coordinate ϵ:

$$g^{(2)}_{\mathrm{out},j} = \frac{\left\langle \int_0^1 \mathrm{d}\epsilon' \ \int_0^1 \mathrm{d}\epsilon'' \ \sigma^\dagger_{\mathrm{out},j}(\epsilon')\sigma^\dagger_{\mathrm{out},j}(\epsilon'')\sigma_{\mathrm{out},j}(\epsilon'')\sigma_{\mathrm{out},j}(\epsilon') \right\rangle}{\left\langle \int_0^1 \mathrm{d}\epsilon' \ \sigma^\dagger_{\mathrm{out},j}(\epsilon')\sigma_{\mathrm{out},j}(\epsilon') \right\rangle^2}, \tag{3.86}$$

The denominator here is simply $N^2_{\mathrm{out},j}$.

The expression in equation 3.86 gates the correlation function across the whole of the retrieval time-bin j. This is analogous to how detector jitter and finite gating times for realistic detectors might be handled [152] by including some detection window function – rather than considering the correlation as a function of the fields at different times or delays. More explictly, the second order correlation function of a quantised field $\hat{E}$ may be written

$$\begin{aligned} g^{(2)}(t_1,t_2) &\sim \left\langle \hat{E}^{(-)}(t_1)\hat{E}^{(-)}(t_2)\hat{E}^{(+)}(t_2)\hat{E}^{(+)}(t_1) \right\rangle \\ &\to \iint \mathrm{d}t' \, \mathrm{d}t'' \ T(t')T(t'') \left\langle \hat{E}^{(-)}(t')\hat{E}^{(-)}(t'')\hat{E}^{(+)}(t'')\hat{E}^{(+)}(t') \right\rangle, \end{aligned} \tag{3.87}$$

where the detector response $T(t)$ weights the correlation function within the detection window. If this window is flat across the length of the pulse and $T(t) \to$ const. – or, as in this case, we integrate over the entire retrieval time-bin – this reduces to the form given in equation 3.86.

If we rewrite equation 3.56 as

$$\begin{aligned} \sigma_{\mathrm{out},j}(\epsilon) &= \int_0^1 \mathrm{d}\epsilon' \left[M^{[j]}_{\mathrm{Ram}}(\epsilon,\epsilon')\sigma_{\mathrm{in},0}(\epsilon') + \sum_{j'=0}^{j} \left(M^{[j']}_{\mathrm{FWM}}(\epsilon,\epsilon')\alpha^\dagger_{\mathrm{in},j-j'}(\epsilon') \right) \right] \\ &\to \tilde{\sigma}_j(\epsilon) + \sum_{j'=0}^{j} \tilde{\alpha}^\dagger_{j'}(\epsilon), \end{aligned} \tag{3.88}$$

then writing out in full to evaluate the numerator of 3.86

$$\left\langle \iint \mathrm{d}\epsilon' \mathrm{d}\epsilon'' \left(\tilde{\sigma}_j(\epsilon') + \Sigma_{j'}\tilde{\alpha}^\dagger_{j'}(\epsilon') \right)^\dagger \left(\tilde{\sigma}_j(\epsilon'') + \Sigma_{j'}\tilde{\alpha}^\dagger_{j'}(\epsilon'') \right)^\dagger \left(\tilde{\sigma}_j(\epsilon'') + \Sigma_{j'}\tilde{\alpha}^\dagger_{j'}(\epsilon'') \right) \left(\tilde{\sigma}_j(\epsilon') + \Sigma_{j'}\tilde{\alpha}^\dagger_{j'}(\epsilon') \right) \right\rangle. \tag{3.89}$$

Again we may see that we can only have non-zero contributions from the terms of the form:

1. $\sim \langle \sigma^\dagger \sigma^\dagger \sigma \sigma \rangle$;
2. $\sim \sum_{k'} \Big[\big(\langle \sigma^\dagger \alpha_{k'} \alpha_{k'}^\dagger \sigma \rangle + \langle \alpha_{k'} \sigma^\dagger \sigma \alpha_{k'}^\dagger \rangle \big) + \big(\langle \sigma^\dagger \alpha_{k'} \sigma \alpha_{k'}^\dagger \rangle + \langle \alpha_{k'} \sigma^\dagger \alpha_{k'}^\dagger \sigma \rangle \big) \Big]$;
3. $\sim \sum_{k,k'} \Big[\langle \alpha_k \alpha_{k'} \alpha_k^\dagger \alpha_{k'}^\dagger \rangle \Big]$.

We look first at the leading term after expanding out:

$$\begin{aligned}
&\left\langle \iint \mathrm{d}\epsilon' \, \mathrm{d}\epsilon'' \; \tilde{\sigma}^\dagger(\epsilon') \tilde{\sigma}^\dagger(\epsilon'') \tilde{\sigma}(\epsilon'') \tilde{\sigma}(\epsilon') \right\rangle \\
&= \int \cdots \int \{\mathrm{d}\epsilon_k\}_{k=0}^{3} \, \mathrm{d}\epsilon' \, \mathrm{d}\epsilon'' \; M^*(\epsilon', \epsilon_0) M^*(\epsilon'', \epsilon_1) M(\epsilon'', \epsilon_2) M(\epsilon', \epsilon_3) \left\langle \sigma^\dagger(\epsilon_0) \sigma^\dagger(\epsilon_1) \sigma(\epsilon_2) \sigma(\epsilon_3) \right\rangle \\
&\to \mathrm{Tr}_\epsilon \Big[\boldsymbol{M}_{\mathrm{Ram}}^{[1]} \vec{\psi}_{\mathrm{in}} \vec{\psi}_{\mathrm{in}}^\dagger \boldsymbol{M}_{\mathrm{Ram}}^{[1]\,\dagger} \Big]^2 \times N_{\mathrm{in},0}^2 g_{\mathrm{in}}^{(2)},
\end{aligned}$$

where we have dropped the subscripts for readability, and denoted the multiple integrals $\int \cdots \int \{\mathrm{d}x_i\}_{i=0}^{N}$, over the set of integration variables x_0, ..., x_N.

Likewise, considering the cross–term contributions to 3.89,

$$\begin{aligned}
&\int \cdots \int \{\mathrm{d}\epsilon_m\}_{m=0}^{3} \, \mathrm{d}\epsilon' \, \mathrm{d}\epsilon'' \; \sum_{k,l} \Big[M_{\mathrm{Ram}}^{[1]\,*}(\epsilon', \epsilon_0) M_{\mathrm{FWM}}^{[k]\,*}(\epsilon'', \epsilon_1) M_{\mathrm{FWM}}^{[l]}(\epsilon'', \epsilon_2) M_{\mathrm{Ram}}^{[1]}(\epsilon', \epsilon_3) \\
&\qquad \times \left\langle \sigma^\dagger(\epsilon_0) \alpha_k(\epsilon_1) \alpha_l^\dagger(\epsilon_2) \sigma(\epsilon_3) \right\rangle \Big] \\
&= \int \cdots \int \{\mathrm{d}\epsilon_m\}_{m=0}^{3} \, \mathrm{d}\epsilon' \, \mathrm{d}\epsilon'' \; \sum_{k,l} \Big[M_{\mathrm{Ram}}^{[1]\,*}(\epsilon', \epsilon_0) M_{\mathrm{FWM}}^{[k]\,*}(\epsilon'', \epsilon_1) M_{\mathrm{FWM}}^{[l]}(\epsilon'', \epsilon_2) M_{\mathrm{Ram}}^{[1]}(\epsilon', \epsilon_3) \\
&\qquad \times \{ N_{\mathrm{in},0} \psi^*(\epsilon_0) \psi(\epsilon_3) \delta(\epsilon_3 - \epsilon_0) \times \delta_{k,l} \delta(\epsilon_2 - \epsilon_1) \} \Big] \\
&= \sum_k N_{\mathrm{in},0} \mathrm{Tr}_\epsilon \Big[\boldsymbol{M}_{\mathrm{Ram}}^{[1]} \vec{\psi}_{\mathrm{in}} \vec{\psi}_{\mathrm{in}}^\dagger \boldsymbol{M}_{\mathrm{Ram}}^{[1]\,\dagger} \Big] \times \mathrm{Tr}_\epsilon \Big[\boldsymbol{M}_{\mathrm{FWM}}^{[k]} \boldsymbol{M}_{\mathrm{FWM}}^{[k]\,\dagger} \Big],
\end{aligned}$$

and by inspection the second symmetric term $\sim \langle \alpha_{k'} \sigma^\dagger \sigma \alpha_{k'}^\dagger \rangle$, is also separable, and yields the same. Likewise, evaluating the antisymmetric terms we find:

$$\begin{aligned}
&\int \cdots \int \{\mathrm{d}\epsilon_m\}_{m=0}^{3} \, \mathrm{d}\epsilon' \, \mathrm{d}\epsilon'' \; \sum_{k,l} \Big[M_{\mathrm{Ram}}^{[1]\,*}(\epsilon', \epsilon_0) M_{\mathrm{FWM}}^{[k]\,*}(\epsilon'', \epsilon_1) M_{\mathrm{Ram}}^{[1]}(\epsilon'', \epsilon_2) M_{\mathrm{FWM}}^{[l]}(\epsilon', \epsilon_3) \\
&\qquad \times \left\langle \sigma^\dagger(\epsilon_0) \alpha_k(\epsilon_1) \sigma(\epsilon_2) \alpha_l^\dagger(\epsilon_3) \right\rangle \Big] \\
&= \int \cdots \int \{\mathrm{d}\epsilon_m\}_{m=0}^{3} \, \mathrm{d}\epsilon' \, \mathrm{d}\epsilon'' \; \sum_{k,l} \Big[M_{\mathrm{Ram}}^{[1]\,*}(\epsilon', \epsilon_0) M_{\mathrm{FWM}}^{[k]\,*}(\epsilon'', \epsilon_1) M_{\mathrm{Ram}}^{[1]}(\epsilon'', \epsilon_2) M_{\mathrm{FWM}}^{[l]}(\epsilon', \epsilon_3) \\
&\qquad \times \{ N_{\mathrm{in},0} \psi^*(\epsilon_0) \psi(\epsilon_2) \delta(\epsilon_2 - \epsilon_0) \times \delta_{k,l} \delta(\epsilon_3 - \epsilon_1) \} \Big] \\
&= \sum_k N_{\mathrm{in},0} \mathrm{Tr}_\epsilon \Big[\boldsymbol{M}_{\mathrm{Ram}}^{[1]} \vec{\psi}_{\mathrm{in}} \vec{\psi}_{\mathrm{in}}^\dagger \boldsymbol{M}_{\mathrm{Ram}}^{[1]\,\dagger} \boldsymbol{M}_{\mathrm{FWM}}^{[k]} \boldsymbol{M}_{\mathrm{FWM}}^{[k]\,\dagger} \Big].
\end{aligned}$$

Lastly, exploiting the commutation relations of the quasi-Bosonic operators on the anti–Stokes mode $\left[\alpha_{[i]}, \, \alpha_{[j]}^\dagger \right] = \delta_{[i],[j]}$, where the bracketed subscript encodes all the relevant operator information.

Acting upon vacuum, and thus discarding normally ordered terms, we are left with only

$$\left\langle \alpha_{[i]}\alpha_{[j]}\alpha^{\dagger}_{[k]}\alpha^{\dagger}_{[l]} \right\rangle = \delta_{[i],[l]}\delta_{[j],[k]} + \delta_{[i],[j]}\delta_{[k],[l]},$$

so putting this in:

$$\begin{aligned}
&\int\cdots\int\{\mathrm{d}\epsilon_m\}_{m=0}^{3}\,\mathrm{d}\epsilon'\,\mathrm{d}\epsilon''\sum_{i,j,k,l}\Big[M^{[i]\,*}_{\mathrm{FWM}}(\epsilon',\epsilon_0)M^{[j]\,*}_{\mathrm{FWM}}(\epsilon'',\epsilon_1)M^{[k]}_{\mathrm{FWM}}(\epsilon'',\epsilon_2)M^{[l]}_{\mathrm{FWM}}(\epsilon',\epsilon_3)\\
&\qquad\qquad\times\left\langle \alpha_i(\epsilon_0)\alpha_j(\epsilon_1)\alpha^{\dagger}_k(\epsilon_2)\alpha^{\dagger}_l(\epsilon_3)\right\rangle\Big]\\
&=\int\cdots\int\{\mathrm{d}\epsilon_m\}_{m=0}^{3}\,\mathrm{d}\epsilon'\,\mathrm{d}\epsilon''\sum_{i,j,k,l}\Big[M^{[i]\,*}_{\mathrm{FWM}}(\epsilon',\epsilon_0)M^{[j]\,*}_{\mathrm{FWM}}(\epsilon'',\epsilon_1)M^{[k]}_{\mathrm{FWM}}(\epsilon'',\epsilon_2)M^{[l]}_{\mathrm{FWM}}(\epsilon',\epsilon_3)\\
&\qquad\qquad\times\{\delta_{i,l}\delta(\epsilon_3-\epsilon_0)\times\delta_{j,k}\delta(\epsilon_2-\epsilon_1)+\delta_{i,j}\delta(\epsilon_1-\epsilon_0)\times\delta_{k,l}\delta(\epsilon_3-\epsilon_2)\}\Big]\\
&=\sum_{k,k'}\left\{\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[k]}_{\mathrm{FWM}}\boldsymbol{M}^{[k]\,\dagger}_{\mathrm{FWM}}\right]\times\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[k']}_{\mathrm{FWM}}\boldsymbol{M}^{[k']\,\dagger}_{\mathrm{FWM}}\right]+\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[k]}_{\mathrm{FWM}}\boldsymbol{M}^{[k]\,\dagger}_{\mathrm{FWM}}\boldsymbol{M}^{[k']}_{\mathrm{FWM}}\boldsymbol{M}^{[k']\,\dagger}_{\mathrm{FWM}}\right]\right\}.
\end{aligned}$$

Note that with the double sum, and with cyclic permutations of the trace that each term arises twice.

And, eventually, recalling the definition of $\boldsymbol{P}_j$ in equation 3.71, we can write:

$$\begin{aligned}
g^{(2)}_{\mathrm{out},j}N^2_{\mathrm{out},j} = g^{(2)}_{\mathrm{in},0}N^2_{\mathrm{in},0}\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\dagger}_{\mathrm{Ram}}\right]^2\\
+\left\{\mathrm{Tr}_\epsilon[\boldsymbol{P}_j]^2-N^2_{\mathrm{in},0}\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\dagger}_{\mathrm{Ram}}\right]^2\right\}\\
+\left\{\mathrm{Tr}_\epsilon\left[\boldsymbol{P}^2_j\right]-N^2_{\mathrm{in},0}\mathrm{Tr}_\epsilon\left[\left(\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\dagger}_{\mathrm{Ram}}\right)^2\right]\right\}.
\end{aligned}$$

To proceed, recall that that $\mathrm{Tr}_\epsilon[\boldsymbol{P}_j]^2 = N^2_{\mathrm{out},j}$, we may write

$$\begin{aligned}
\left(g^{(2)}_{\mathrm{out},j}-1\right)N^2_{\mathrm{out},j} = \left(g^{(2)}_{\mathrm{in},0}-1\right)N^2_{\mathrm{in},0}\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\dagger}_{\mathrm{Ram}}\right]^2\\
+\left\{\mathrm{Tr}_\epsilon\left[\boldsymbol{P}^2_j\right]-N^2_{\mathrm{in},0}\mathrm{Tr}_\epsilon\left[\left(\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\dagger}_{\mathrm{Ram}}\right)^2\right]\right\},
\end{aligned}\tag{3.90}$$

The final step is now to evaluate the form of $\mathrm{Tr}_\epsilon\left[\boldsymbol{P}^2_j\right]$.

To begin, we first observe that

$$\mathrm{Tr}_\epsilon\left[\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\dagger}_{\mathrm{Ram}}\right]^2 = \mathrm{Tr}_\epsilon\left[\left(\boldsymbol{M}^{[1]}_{\mathrm{Ram}}\vec{\psi_{\mathrm{in}}}\cdot\vec{\psi_{\mathrm{in}}}^{\,\dagger}\boldsymbol{M}^{[1]\dagger}_{\mathrm{Ram}}\right)^2\right] = |\chi C_{\mathrm{s}}|^4|C_{\mathrm{s}}|^4\,E^4\,|\mu|^4.$$

Moving forward, to simplify the following integrals to evaluate the statistics in the first retrieval bin, we recall the definitions for the kernels 3.49, 3.53 and 3.54, and define

$$\begin{aligned}
\boldsymbol{K}_0 &= \mathrm{e}^{h}\boldsymbol{M}\vec{\psi}_{\mathrm{in}}\vec{\psi}^{\dagger}_{\mathrm{in}}\boldsymbol{M}^{\dagger}\mathrm{e}^{h^*},\\
\boldsymbol{K}_1 &= \quad\mathrm{e}^{h}\boldsymbol{M}\boldsymbol{M}^{\dagger}\mathrm{e}^{h^*},\\
\boldsymbol{K}_2 &= \qquad\boldsymbol{M}_{\boldsymbol{c}}\boldsymbol{M}^{\dagger}_{\boldsymbol{c}}.
\end{aligned}$$

With a little working we find

$$[\boldsymbol{K}_0]_{i,j} = E\,|\mu|^2 \times \mathrm{e}^{h\epsilon_i + h^*\epsilon_j} \tag{3.91}$$

$$[\boldsymbol{K}_1]_{i,j} = E \times \mathrm{e}^{h\epsilon_i + h^*\epsilon_j} \tag{3.92}$$

$$[\boldsymbol{K}_2]_{i,j} = \mathrm{e}^{h\epsilon_i + h^*\epsilon_j}\frac{1}{\xi}\left[\mathrm{e}^{\xi\min(\epsilon_i,\epsilon_j)} - 1\right], \tag{3.93}$$

where to arrive at the last line, we have used the Heaviside step–function to set the integral limits:

$$\begin{aligned}[\boldsymbol{K}_2]_{i,j} &= \int_0^1 \mathrm{d}\epsilon'\ \Theta(\epsilon_i - \epsilon')\mathrm{e}^{h(\epsilon_i-\epsilon')}\Theta(\epsilon_j - \epsilon')\mathrm{e}^{h^*(\epsilon_j-\epsilon')} \\ &= \mathrm{e}^{h\epsilon_i + h^*\epsilon_j}\int_0^{\min(\epsilon_i,\epsilon_j)} \mathrm{d}\epsilon'\ \mathrm{e}^{+\xi\epsilon'} \ = \ \mathrm{e}^{h\epsilon_i + h^*\epsilon_j}\frac{1}{\xi}\left[\mathrm{e}^{\xi\min(\epsilon_i,\epsilon_j)} - 1\right].\end{aligned}$$

So, as previously we can consider the terms which will arise – noting that the trace is invariant under cyclic permutations – firstly:

$$\begin{aligned}\mathrm{Tr}_\epsilon[\boldsymbol{K}_0\boldsymbol{K}_1] \propto \mathrm{Tr}_\epsilon[\boldsymbol{K}_1\boldsymbol{K}_1] &\propto \int_0^1 \mathrm{d}\epsilon'\ \int_0^1 \mathrm{d}\epsilon''\ \mathrm{e}^{h\epsilon' + h^*\epsilon''}\mathrm{e}^{h\epsilon'' + h^*\epsilon''} \\ &= \left(\int_0^1 \mathrm{d}\epsilon'\ \mathrm{e}^{-\xi\epsilon'}\right)\left(\int_0^1 \mathrm{d}\epsilon''\ \mathrm{e}^{-\xi\epsilon''}\right) \\ &= E^2,\end{aligned}$$

and gathering up the prefactors:

$$\mathrm{Tr}_\epsilon[\boldsymbol{K}_0\boldsymbol{K}_1] = E^4|\mu|^2 \qquad \text{and} \qquad \mathrm{Tr}_\epsilon[\boldsymbol{K}_1\boldsymbol{K}_1] = E^4. \tag{3.94}$$

Then for the causal Green's functions

$$\begin{aligned}\mathrm{Tr}_\epsilon[\boldsymbol{K}_0\boldsymbol{K}_2] \propto \mathrm{Tr}_\epsilon[\boldsymbol{K}_1\boldsymbol{K}_2] &\propto \int_0^1 \mathrm{d}\epsilon'\ \int_0^1 \mathrm{d}\epsilon''\ \mathrm{e}^{-\xi(\epsilon'+\epsilon'')}\left[\mathrm{e}^{\xi\min(\epsilon_i,\epsilon_j)} - 1\right] \\ &= \int_0^1 \mathrm{d}\epsilon'\ \left\{\int_0^{\epsilon'} \mathrm{d}\epsilon''\ \mathrm{e}^{-\xi\epsilon'} + \int_{\epsilon'}^1 \mathrm{d}\epsilon''\ \mathrm{e}^{-\xi\epsilon''}\right\} - E^2 \\ &= \frac{2\mathrm{e}^{-\xi}}{\xi^2}\Big(\sinh(\xi) - \xi\Big) = \frac{\Sigma}{\xi^2},\end{aligned}$$

where we have defined $\Sigma = 2\mathrm{e}^{-\xi}[\sinh(\xi) - \xi]$. So again, putting in the prefactors

$$\mathrm{Tr}_\epsilon[\boldsymbol{K}_0\boldsymbol{K}_2] = \frac{E|\mu|^2\Sigma}{\xi^3} \qquad \text{and} \qquad \mathrm{Tr}_\epsilon[\boldsymbol{K}_1\boldsymbol{K}_2] = \frac{E\Sigma}{\xi^3}. \tag{3.95}$$

And, finally:

$$\begin{aligned}\mathrm{Tr}_\epsilon[\boldsymbol{K}_2\boldsymbol{K}_2] &= \frac{1}{\xi^2}\int_0^1 \mathrm{d}\epsilon'\ \int_0^1 \mathrm{d}\epsilon''\ \mathrm{e}^{-\xi(\epsilon'+\epsilon'')}\left[1 - 2\mathrm{e}^{\xi\min(\epsilon',\epsilon'')} + \mathrm{e}^{2\xi\min(\epsilon',\epsilon')}\right] \\ &= \frac{1}{\xi^2}\Big[E^2 - 2\left(\frac{\Sigma}{\xi^2} - E^2\right) + \int_0^1 \mathrm{d}\epsilon'\ \left\{\int_0^{\epsilon'} \mathrm{d}\epsilon''\ \left(\mathrm{e}^{-\xi\epsilon'+\xi\epsilon''}\right) + \int_{\epsilon'}^1 \mathrm{d}\epsilon''\ \left(\mathrm{e}^{+\xi\epsilon'-\xi\epsilon''}\right)\right\}\Big] \\ &= \frac{1}{\xi^2}\Big[-\left(E^2 + \frac{2\Sigma}{\xi^2}\right) + 2\left(\frac{1-E}{\xi}\right)\Big].\end{aligned} \tag{3.96}$$

Summing up contributions

$$\begin{aligned}\mathrm{Tr}_\epsilon\left[\boldsymbol{P}_1^2\right] &= |\chi C_\mathrm{s}|^4 \Big\{ |C_\mathrm{s}|^4 N_{\mathrm{in},0}^2 \times \mathrm{Tr}_\epsilon[\boldsymbol{K}_0\boldsymbol{K}_0] \\ &\qquad + |C_\mathrm{s}C_\mathrm{a}x|^2 N_{\mathrm{in},0} \times \mathrm{Tr}_\epsilon[\boldsymbol{K}_0\boldsymbol{K}_1 + \boldsymbol{K}_1\boldsymbol{K}_0 + \boldsymbol{K}_0\boldsymbol{K}_2 + \boldsymbol{K}_2\boldsymbol{K}_0] \\ &\qquad + |C_\mathrm{a}x|^4 \times \mathrm{Tr}_\epsilon[\boldsymbol{K}_1\boldsymbol{K}_1 + \boldsymbol{K}_1\boldsymbol{K}_2 + \boldsymbol{K}_2\boldsymbol{K}_1 + \boldsymbol{K}_2\boldsymbol{K}_2] \Big\} \\ &= |\chi C_\mathrm{s}|^4 \left\{ |C_\mathrm{s}|^4 N_{\mathrm{in},0}^2 E^4 |\mu|^4 + |C_\mathrm{s}C_\mathrm{a}x|^2 N_{\mathrm{in},0} 2E|\mu|^2 G_0(\xi) + |C_\mathrm{a}x|^4 G_1(\xi) \right\}, \end{aligned} \tag{3.97}$$

we can finally write

$$\begin{aligned}\left(g_{\mathrm{out},1}^{(2)} - 1\right) N_{\mathrm{out},1}^2 = |\chi C_\mathrm{s}|^4 \Big\{ &\left(g_{\mathrm{in},0}^{(2)} - 1\right) N_{\mathrm{in},0}^2 \times |C_\mathrm{s}|^4 E^4 |\mu|^4 \\ &+ |C_\mathrm{a}x|^2 \left[2E|\mu|^2 N_{\mathrm{in},0} |C_\mathrm{s}|^2 G_0(\xi) + |C_\mathrm{a}x|^2 G_1(\xi) \right] \Big\}, \end{aligned} \tag{3.98}$$

where we have defined the (somewhat cumbersome) functions

$$G_0(\xi) = E^3 + \frac{\Sigma}{\xi^3} \tag{3.99}$$

$$G_1(\xi) = E^4 + \frac{2\Sigma}{\xi^4}\left(\xi E - 1\right) + 2\left(\frac{1-E}{\xi^3}\right) - \left(\frac{E}{\xi}\right)^2. \tag{3.100}$$

Again, inspecting the expression for the $g^{(2)}$ term-by-term

$$\left(g_{\mathrm{out},1}^{(2)} - 1\right) N_{\mathrm{out},1}^2 = |\chi C_\mathrm{s}|^4 \left\{ \alpha |C_\mathrm{s}|^4 \times \left(g_{\mathrm{in},0}^{(2)} - 1\right) N_{\mathrm{in},0}^2 + \beta\, |C_\mathrm{a}x|^2 |C_\mathrm{s}|^2 N_{\mathrm{in},0} + \gamma\, |C_\mathrm{a}x|^4 \right\}, \tag{3.101}$$

where the overall prefactor $\sim |C_\mathrm{s}|$ arises since we are considering the statistics of light retrieved into the signal mode. Since the $g^{(2)}$ is a second order correlation metric, it depends on two photons in the retrieval mode. Within this model,[9] these two photons may (a) both originate from the input signal $\sim |C_\mathrm{s}|^4$; (b) both from FWM $\sim |C_\mathrm{a}x|^4$; or (c) one from each $\sim |C_\mathrm{s}|^2 |C_\mathrm{a}x|^2$. The coefficients α, β, γ are given below.

$\alpha = E^4\|\mu\|^4$	target input–output mapping of the pure Raman process, from $E^2\|\mu\|^2 = \mathrm{Tr}_\epsilon\left[\boldsymbol{M}\vec{\psi_\mathrm{in}} \cdot \vec{\psi_\mathrm{in}}^\dagger \boldsymbol{M}^\dagger\right]$
$\beta = 2E\|\mu\|^2 G_0(\xi)$	cross term involving both Stokes/anti-Stokes modes. The two components of G_0 correspond to when the anti-Stokes scattering event occured in the read-in (E^3) or read-out ($\sim \mathrm{e}^{-\xi}[\sinh(\xi) - \xi]$) time bins.
$\gamma = G_1(\xi)$	Spontaneous FWM which builds up in the first two time bins.

3.2.3 Theory Summary

We now have a metric to predict the second order statistics of the memory action, in terms of the statistics of the input state. If we consider a single photon incident upon the memory, the Fock state $|1\rangle$ with $g_\mathrm{in}^{(2)} = 0$, the statistics of the retrieved state $g_{\mathrm{out},j}^{(2)}$ suffers pollution via the additional terms in equation 3.98. By choosing the cavity parameters, we have experimental control over the suppression

[9] Neglecting, for instance, spontaneous Raman scattering into the retrieval mode from imperfect pumping.

Variable (eq.)	**Strong coupling** ($\xi \gg 1$)	**Weak coupling** ($\xi \ll 1$)
E (3.75)	$1/\xi$	$1 - \xi/2$
g (3.82)	$1/\xi$	$3/2 - 7\xi/6$
Σ (3.95)	1	$\xi^3/3$
G_0 (3.99)	$2/\xi^3$	$4/3 - 11\xi/6$
G_1 (3.100)	$2/\xi^3$	$11/6 - 47\xi/15$

Table 3.2: Limiting behaviour of parameters, to leading order in terms of ξ, defined in equation 3.74

factor x, which we seek to minimise, and remove the four–wave mixing noise.

Before proceeding to the experimental implementation, we identify a few implications of these results in some limiting cases.

With our expression for the $g^{(2)}$, for which, in general we might be sending in single photon Fock states, i.e. with $g^{(2)}_{\text{in},0} = 0$, and where $N_{\text{in},0} \leq 1$ may be interpreted as the heralding efficiency of the source. After some rearranging we find

$$g^{(2)}_{\text{out},1} = \frac{2N_{\text{in},0}\,|C_\text{s}C_\text{a}x|^2\,E|\mu|^2\,[Eg + G_0] + |C_\text{a}x|^2\,\left[g^2 + G_1\right]}{N_{\text{in},0}\,|C_\text{s}|\,E^2|\mu|^2 + |C_\text{a}x|^2\,g}. \tag{3.102}$$

Here we recall the definitions

$$E^2|\mu|^2 = \text{Tr}_\epsilon\left[\boldsymbol{M}\vec{\psi_\text{in}}\cdot\vec{\psi_\text{in}}^{\,\dagger}\boldsymbol{M}^\dagger\right]$$
$$g = \text{Tr}_\epsilon\left[\boldsymbol{M_c}\boldsymbol{M_c}^\dagger\right] + \text{e}^{h+h^*}\,\text{Tr}_\epsilon\left[\boldsymbol{M}\boldsymbol{M}^\dagger\right]$$

and G_0 (G_1) are the contribution to correlations in the output signal mode which involve one (two) FWM event(s).

We observe that in the case of strong suppression, that is for $|x|^2 \ll N_{\text{in},1}$,

$$g^{(2)}_{\text{out},1} \approx 2\left[1 + \frac{G_0}{Eg}\right]\frac{1}{\text{SNR}}. \tag{3.103}$$

The overall coupling strength is parameterised by $\xi = -\mathfrak{Re}\{h\}$, defined in equation 3.74, since h provides the scaling for the time evolution of the Green's function 3.43. In the absence of FWM and Stark shift $h = C_\text{s}$ is simply the cavity enhanced Raman coupling. This coupling depends on the signal frequency – via detuning from the atoms and the absorption/dispersion it experiences within the cavity, the Rabi frequency, and the optical depth of the ensemble.

The behaviour of some of the different parameters defined in the preceeding sections are shown in table 3.2 for the cases of weak coupling $\xi \ll 1$ and strong coupling $\xi \gg 1$, and we discuss the resulting performance parameters in the following sections.

3.2.3(a) Memory Efficiency

The memory efficiency achievable is strongly dependent on intracavity loss. We find the two limiting cases

$$\lim_{\xi\to 0}(\eta_{\text{mem}}) = (1-\xi)\left|\chi C_{\text{s}}^2\mu\right|^2, \tag{3.104}$$

$$(\eta_{\text{mem}})_{\xi\gg 1} \to \left|\frac{\chi C_{\text{s}}^2\mu}{\xi}\right|^2. \tag{3.105}$$

Thus at low coupling the efficiency rises as the coupling is increased, say via W or d. At high coupling, this result is simplified if the we have strong noise suppression, and we are far detuned from resonance, where from equation 3.74 we see that $\xi \to -2C_{\text{s}}$, and $\eta_{\text{mem}} \to |\chi\mu|^2/4$ which defines the maximal achievable efficiency in the cavity Raman memory, determined by the mode overlap μ and the roundtrip transmission χ.

3.2.3(b) Statistics

First we consider limiting cases of the signal to noise ratio, where we recall

$$\text{SNR} = N_{\text{in},0}|\mu|^2\left|\frac{C_{\text{s}}}{C_{\text{a}}x}\right|^2 \times \frac{E^2}{g} \tag{3.106}$$

$$\text{where :} \qquad \frac{E^2}{g} \to \begin{cases} 1/\xi & \text{strong coupling;} \\ 2/3 - 4\xi/27 & \text{weak coupling.} \end{cases} \tag{3.107}$$

And similarly considering the dynamical components governing the $g^{(2)}$ in the *strong suppression regime* given by the approximation 3.103

$$\left[1+\frac{G_0}{Eg}\right] \to \begin{cases} 1+2/\xi & \text{strong coupling;} \\ 17/9 - 7\xi/81 & \text{weak coupling.} \end{cases} \tag{3.108}$$

With this in mind we see that *in this limit* from the approximation 3.103

$$g^{(2)}_{\text{out},1} \sim \begin{cases} 2\xi \times \left[N_{\text{in},0}|\mu|^2\right]^{-1}\left|\frac{C_{\text{a}}x}{C_{\text{s}}}\right|^2 + \mathcal{O}\left(\frac{1}{\xi}\right) & \xi \gg 1; \\ 2\left(\frac{17}{3}+\xi\right) \times \left[N_{\text{in},0}|\mu|^2\right]^{-1}\left|\frac{C_{\text{a}}x}{C_{\text{s}}}\right|^2 + \mathcal{O}\left(\xi^2\right) & \xi \ll 1. \end{cases} \tag{3.109}$$

So in the strong coupling case, the output $g^{(2)}$ rises linearly with ξ thus our condition for suppression of FWM noise is that $|x|^2 \ll N_{\text{in},0}/\xi$.

With the efficiency and autocorrelation statistics of the memory investigated in these limits, we conclude that for optimal operation we seek to work with the *lowest* coupling ξ for which the efficiency saturates.

3.2.3(c) Optimal Mode

We highlighted that the interaction within this model is *exactly* single mode, in that the temporal mode selected by the memory is given by equation 3.55. For a given control pulse, which defines the

time coordinate ϵ, if one prepares the input mode:

$$\sigma_{\text{opt}}(\epsilon) = \left(\mathrm{e}^{\xi} E\right)^{-1/2} \mathrm{e}^{-h^*\epsilon} \tag{3.110}$$

$$\Rightarrow \quad s_{\text{opt}}(t) \propto \frac{\Omega(t)}{\sqrt{W}} \exp\left[-h^* \int_0^t \mathrm{d}t' \left|\frac{\Omega(t')}{\sqrt{W}}\right|^2\right], \tag{3.111}$$

results in $\mu = 1$. In the second line, we have recalled the normalisation employed in equation 3.25 to transform back to the laboratory time frame.

In the case of low coupling $|h| \ll 1$, we see that the signal profile follows the temporal profile of the control field. However, as ξ is increased, the mode addressed by the memory becomes distorted in amplitude and phase.

Alternatively, more practically, the above relationship may be inverted to show how, for a given W, the control field must be modified in order to address some arbitrary target signal. This is applicable *independently* in both storage and retrieval interactions [138], such that the interaction remains operating on a chosen basis of signal temporal modes. This will be discussed further in the second part of this thesis.

3.3 Cavity Engineering

Having developed a theoretical description of the Raman memory in a low finesse cavity, and identified some of the figures of merit and critical parameters for memory performance and noise suppression, we now move on to present the first experimental demonstration of a single-photon-level Raman memory in a cavity.

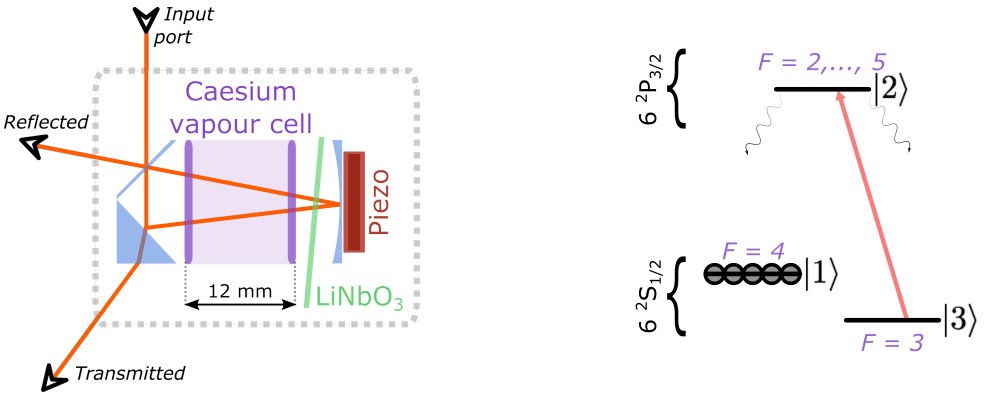
For this proof of principle experiment, we choose to implement a bulk optics design. If one considers meeting the resonance/anti-resonance condition between signal and anti-Stokes, we seek the condition $(M + 1/2)\text{FSR} = 2\Delta_0$ where we call the integer M the "cavity order".[10] For a given bandwidth, one observes that the visibility – which relates to the suppression factor x introduced previously – is enhanced for lower M; however, as the FSR is inversely proportional to the length this requires a smaller cavity, thus for a table top demonstration we take $M = 2$ requiring an intracavity roundtrip path length of ~ 40 mm (before any atoms are included).

The original design for this experment was a dual-port ring cavity (figure 3.3), instead of the single-port (figure 3.2) configuration considered in the previous section. The aim of this being to achieve critical coupling, such that the incident field is fully coupled into the cavity at the input coupler, and giving greater optical access; however, we note that this increases the intracavity loss and thereby reduces the suppression available, and a photon retrieved from the memory inside the cavity has no external field to interfere with and could exit via either port.

3.3.1 Cavity Memory Experiment

The system analysed in the present work is comprised of a ring cavity, passing through caesium-filled vapour cell twice each roundtrip as illustrated in figure 3.3a. The vapour cell (Precision Glassblowing)

[10] This terminology was introduced in [153] whereby M defines the number of cavity modes that lie in between the Stokes and anti-Stokes frequencies in figure 3.1. In other words $(M + 1/2)\text{FSR}_M^{(0)} \approx 2\Delta_s$, where $\text{FSR}_M^{(0)}$ is the free spectral range of an M–order cavity far from the atomic resonances. Note that meeting the exact resonance condition requires evaluation of the modification of the cavity spectrum accounting for the complex susceptibility of the atoms [154], discussed in the following section.



(a) (b)

Figure 3.3: (a) Schematic of the ring cavity design experimentally implemented; (b) Energy level Λ configuration of the caesium D_2 transition. The ensemble is initialised into the $F = 4$ hyperfine ground state via optical pumping.

of length 12 mm and diameter 10 mm also contained 10 Torr of neon buffer gas. To control the number density of caesium in the vapour phase, the temperature of the cell was controlled by ohmic heating of a wire wrapped around the cell, with thermistors to monitor the temperature, this is discussed further in section 3.3.1(b). The concave mirror was mounted on a piezo–electric stack such that the cavity length could be actively stabilised, this is discussed further in section 3.3.1(a).

In order to meet the "triple resonance" condition – that is resonance with the signal and control fields and anti-resonance with the anti-Stokes a birefringent crystal (200 μm $LiNbO_3$, double passed) is included in the intracavity path so as to gain some independent control of the relative internal path length for orthogonal polarisations. This was held in a temperature controlled mount, on a rotation stage to allow for angle tuning (for initial alignment) and thermal tuning. The birefringence tuning set up was built by T.F.M.C.

A schematic of the cavity, and the initialised energy configuration of the atomic ensemble are shown in figures 3.3.

The stability of the cavity is critical for these experiments: in particular, the length of the cavity and temperature of the caesium cell were actively stabilised. I briefly discuss the control elements for which I was responsible in the following sections; further details, and some discussion of the experimental evolution, of the mechanical and thermal stabilisation infrastructure are included in the appendix A.1.

3.3.1(a) Mechanical Stabilisation

The cavity length is actively stabilised, by controlling the voltage applied to a piezo-electric actuator (Noliac) on which the curved mirror in the cavity is mounted. In order to do so, we make use of the birefringence of the cavity[11] and generate an error signal based on the Hänsch-Couillaud method [155]. The locking signal arises from the phase difference aquired by orthogonal polarizations in traversing the cavity. By sending in light far off resonant with the atomic caesium, prepared in diagonal polarisation (D) with respect to the cavity, the H and V components will pick up a relative phase between them and thus the reflected light will in general be elliptically polarized. Projecting this onto the horizontal/vertical basis with a PBS, and looking at the difference signal between the

[11] Recall, the intracavity birefringence is introduced to facilitate independent tuning of the cavity response for the signal and control fields.

two ports will yield information about the relative phase, and hence the cavity length. We prepare a HeNe laser (Thorlabs) with diagonal polarisation, counter–propagating with the control and signal, and monitor the difference signal between H and V. The outputs of the two photodiodes are monitored by the Arduino which performs PID logic based on the open source Arduino library [156] to feed back on the piezo driver voltage on the kHz timescale.

For measurements approaching the single photon level, we enter a regime in which detection events are occuring at the Hz timescale. For these measurements, it is therefore vital that the system remains stable for extended periods of time. In order to enable this, bearing in mind the finite resolution of the two Arduino DAC outputs, a coarse grained variable offset term is added to the fine grained PID output and increments automatically as the PID output approaches the edge of its range. For this, the two signals are added in ratio such that the slow (Hz) offset is scaled in order to access the full voltage range of the piezo driver, whilst the PID (kHz) output varies only over a voltage range corresponding to a displacement on the order of a wavelength.

With the addition of the slow feed back the cavity was stable to $\pm 50\,\mathrm{MHz}$ for periods on the order of an hour, in contrast to the PID alone which was only stable for periods of minutes.

3.3.1(b) Thermal Stabilisation

Similarly, the cavity response is strongly dependent on the optical depth of the caesium, and hence temperature stability is also critical. We estimate the temperature control infrastructure of the final version remained stable to $\pm 0.1\,\mathrm{K}$ over the same timescale. In our typical operating conditions, a change in temperature on the order of 1 K can shift the resonance at the signal field on the order of 1/2 the acceptance bandwidth.

The exterior of the cell is heated by driving a DC current through a quad-twisted wire, where the "quad-twist" minimises magnetic fields introduced. wrapped around the cell. Micro thermistors between the wire and the cell are used to monitor the temperature at various points. The resistance of the thermistors are logged by a second Arduino Due and passed to a Labview script which handles a PID feedback system on the heater wire voltage to stabilise the temperature of the cell against temperature drifts in the lab.

At our operating temperatures, caesium is in a liquid phase, which in the cell accumulates at the coldest point. Due to the geometry of the cell, the windows of the cell are the furthest point from the heater wire; and so to prevent caesium build up on the faces of the cell which has previously dramatically reduced the cell transmission due to scattering, a copper cold finger is bonded with GE varnish to the circumference of the cell in the middle of its length. The cold finger is then driven by a constant voltage Peltier element fixed to a heat sink external to the optical table. The two cell faces are additionally covered with copper foil, with a small hole for optical access, which provides thermal contact with the heater wire and cell mount.

3.3.1(c) Measurement of Cavity Response

To characterise the cavity response, scans were conducted with the cavity locked to maximise transmission at a detuning of $\Delta = +6\,\mathrm{GHz}$ from the $|6^2S_{1/2}; F = 3\rangle \rightarrow |6^2P_{3/2}; F = 2\rangle$, $|6^2P_{3/2}; F = 3\rangle$ crossover transition, shown by the circle and vertical line in figure 3.4. I subsequently refer this as *the blue transition* which provides the frequency reference for the laser detuning. A continuous wave diode laser (Toptica DL Pro, <1 MHz linewidth) locked to this resonance both provides a frequency

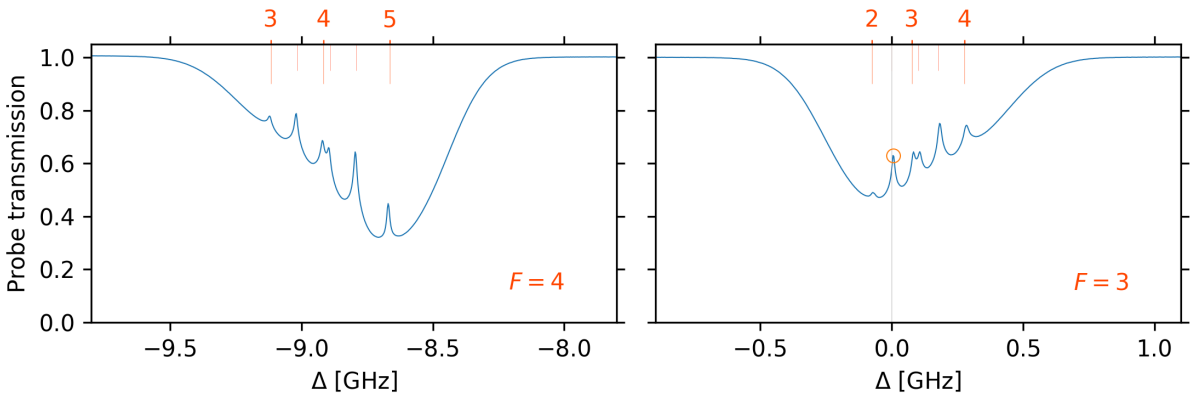


Figure 3.4: Doppler free saturation-absorption spectrum of the caesium D_2 line. This provides the locking signal for the diode laser. The laser detuning, Δ, is measured from the $|6^2S_{1/2}; F = 3\rangle \rightarrow |6^2P_{3/2}; F = 2\rangle$, $|6^2P_{3/2}; F = 3\rangle$ crossover transition.

reference and optically pumps the ensemble into the $F = 4$ hyperfine ground state.[12] A modelocked Ti:Sa (Spectra Physics Tsunami ~ 1 GHz linewidth, 80 MHz repetition rate) was then scanned across both of the D_2 resonances, and the cavity transmission recorded (Fig. 3.8), where the mode-hop free tuning range limited the scans to a range of ~30 GHz. See the later section 3.4.1 for further discussion.

Note that, to fully model the cavity response observed with a given probe laser, the cavity response model must also be convolved with the frequency profile of the probe: which in this case we take to be a Gaussian $g_{\sigma_p}(\omega - \omega_0)$, with FWHM (in intensity) of 0.5 GHz.

3.3.2 Atoms In A Cavity

We briefly take a step back from the experiment and consider the general optical response of atoms in a cavity, and investigate how this modifies the resonance conditions – which are critical for the experiment.

The following analysis was published in [154], and we use this as the basis to show how, in practice, we may meet the resonance conditions necessary in the previous section of this chapter.

We therefore begin our analysis by considering a more simple linear Fabry-Perot type cavity containing some atomic ensemble, depicted schematically in figure 3.5, the transmission through which can be expressed in the familiar form:

$$E_{\text{trans}} = t_{\text{out}} \frac{1}{1-\zeta} \, t_{\text{in}} \, E_{\text{in}} \tag{3.112}$$

where $\zeta = \tau \, r_{\text{in}} \, \tau \, r_{\text{out}}$ is the round trip complex transmission between the outcoupling and incoming mirror, with τ accounting for the single pass loss and phase accumulated passing through the cavity.

To include the contribution to the atomic ensemble on the cavity response, we begin with the analysis of Siddons *et al.* [157] to model the electric susceptibility of the vapour and have based our simulation on the ElecSus computer program [158] published by the same group. This is then incorporated into our own extended framework to describe the entire atom plus cavity system.

In the expression for the cavity transfer function in equation 3.112, we can account for the

[12] This resonance is near the strongest part of the absorption spectrum, and provides the strongest feature to lock to.

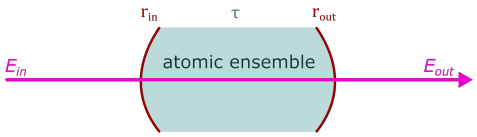


Figure 3.5: We consider the transmission response of a cavity contaning some atomic medium, where intracavity ensemble could be warm vapours, cold atoms or solid state systems.

frequency-dependent dispersion and absorption of the atomic ensemble[13] in the roundtrip transmission term

$$\zeta = |\zeta(\omega)|\mathrm{e}^{\mathrm{i}\phi(\omega)}, \tag{3.113}$$

at laser frequency ω. For the evaluation of the phase accumulation across the cavity, it is necessary to consider the refractive index integrated over the length L of the ensemble, where the refractive index $n = n' + \mathrm{i}n''$ is related to the susceptibility, χ, by

$$n = (1+\chi)^{1/2}. \tag{3.114}$$

Separating real ($'$) and imaginary ($''$) parts, we have

$$n' \approx 1 + \frac{\chi'}{2} \tag{3.115}$$

$$n'' \approx \frac{\chi''}{2}, \tag{3.116}$$

for small χ, as is typically the case in our operating regime.

The intracavity transmission factor, τ, in equation 3.112 can then be written

$$\tau(\omega) = \exp\left[\int_{z=0}^{L} \mathrm{d}z \frac{\omega z}{c}\left\{\frac{-\chi''(\omega,z)}{2} + \mathrm{i}\left(1 + \frac{\chi'(\omega,z)}{2}\right)\right\}\right], \tag{3.117}$$

and can be substituted into ζ.

The susceptibility can be expressed in terms of the contributions of the dipole moments for each of the atomic transitions. For a given dipole transition between states $|i\rangle$ and $|f\rangle$, the contribution to the susceptibility is of the form [159]

$$\chi_{if} = c_{if}^2 \frac{N_i D_{if}^2}{\hbar\varepsilon_0} s_\lambda(\Delta_{if}), \tag{3.118}$$

where c_{if} is the transition strength factor for the given atomic transition, for instance tabulated in [143] for caesium as considered later in the experimental part of this work, $D_{if} = \langle i|e\mathbf{r}|f\rangle$ is the dipole matrix element, given approximately by $D_{if} \approx \langle i||e\mathbf{r}||f\rangle/3$, for linear polarisations [143] assumed in the following analysis, and is the same for all dipole transitions of the D_2 line. N_i is the number density of the atoms in state $|i\rangle$ interacting with the light field, and $s_\lambda(\Delta_{if})$ is a shape factor depending on the detuning, $\Delta_{if} = \omega - \omega_{if}$, from the transition $i \to f$, and some set of physical parameters, λ, as elaborated below.

With the total susceptibility $\chi = \sum_{i,f} \chi_{if}$, we obtain the complex refractive index of the vapour, giving us the dispersion (real part) and absorption (imaginary part) of the vapour.

[13] Additionally, the reflectivities may be polarisation dependent and are not necessarily real scalars, for instance due to stress induced birefringence from vapour cell windows.

For the bare atoms we label the lineshape $f_\gamma(\Delta)$, which is determined solely by Δ and the natural linewidth of the transition, γ (half-width at half-maximum). For a given transition $i \to f$, the shape factor can be derived [144] as

$$\begin{aligned} f_\gamma(\Delta) &= \frac{1}{-\Delta + \mathrm{i}\gamma} + \frac{1}{(\Delta + 2\omega_{if}) + \mathrm{i}\gamma} \\ &\simeq \frac{-\Delta + \mathrm{i}\gamma}{\Delta^2 + \gamma^2}. \end{aligned} \tag{3.119}$$

In the second line we have taken that Δ, $\gamma \ll \omega_{if}$ and made the rotating wave approximation. In practice, $s_\lambda(\Delta)$ must then incorporate the Lorentzian natural linewidth together with contributions from homogeneous and inhomogeneous broadening mechanisms. In atomic vapour systems these may include collisional broadening and Doppler broadening; or in solid state ensembles similar broadening results from phonons and local charge fluctuations or random strain distribution.

Homogeneous broadening contributes an additional Lorentzian component to the natural linewidth by adding an additional decay rate to the natural decay rate of the excited state, identically to all emitters. Here, we consider the specific case of collisional broadening in caesium vapour, where the contribution from pressure broadening and shift of caesium with various buffer gasses is well studied in the literature, see for instance [160, 161]. Nonetheless, analogous broadening occurs in solid state systems. Collisional broadening can therefore be described by including an additional component to the Lorentzian natural linewidth [144] f:

$$\gamma_\mathrm{P} = \gamma + \gamma_\mathrm{coll}, \tag{3.120}$$

where γ is the natural linewidth, and $\gamma_\mathrm{coll} \propto n\bar{v}_\mathrm{th}$ is the contribution from collision induced broadening, where n is the number density and $\bar{v}_\mathrm{th}$ is the thermal velocity of the atoms and the constant of proportionality relates to their collisional cross section. For experiments with relatively low atomic densities of caesium, the pressure broadening is dominated by the buffer gas. Similarly, the buffer gas induces a shift in the energy of the transitions, where the shift is also linear in the pressure.

Inhomogeneous broadening due to the varying local environment of each emitter has the effect of convolving the refractive index over the probability distribution of shifted resonance frequencies. In the case of Doppler broadening this means convolving over the thermal velocity distribution of the atoms, due to the Doppler shifted frequency $\omega' = \omega(1 - v_z/c) = \omega - \delta_v$, seen by an atom with a velocity component v_z parallel to the direction of beam propagation. The resulting shape s, as a function of detuning is then given by the convolution

$$\begin{aligned} s_{\gamma,\sigma}(\Delta) &= [f_\gamma * g_\sigma](\Delta) \\ &= \int_{-\infty}^{\infty} f_\gamma(\Delta - \delta_v) g_\sigma(\delta_v)\, \mathrm{d}\delta_v, \end{aligned} \tag{3.121}$$

where $g_\sigma(\delta_v)$ is the Gaussian distribution of frequency shift due to velocity in one dimension

$$g_\sigma(\delta_v) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-1}{2}\left(\frac{\delta_v}{\sigma}\right)^2\right], \tag{3.122}$$

and

$$\delta_v = \frac{\omega v_z}{c} \tag{3.123}$$

$$\sigma = \frac{\omega}{c}\sqrt{\frac{k_B T}{m}}. \tag{3.124}$$

The resulting shape function is related to the Faddeeva function (if f_γ is taken before making the rotating wave approximation in equation 3.119) [157], and arises in a number of physical systems; nevertheless, it is difficult to handle numerically and has therefore attracted an amount of research over the years into finding either analytic approximations (e.g. [162]) or more recently developing efficient algorithms (e.g. [163, 164]). The absorptive component results in the well known Voigt function. On the other hand, atomic *dispersion* is in general more difficult to measure directly, with experiments requiring actively stabilised interferometry, for instance [165]. Our approach enables access to the dispersive part of the lineshape function through the phase accummulated by the probe light on each round trip of the cavity, which is in turn mapped onto intensity.

Assuming that the natural linewidth γ (and likewise the pressure width γ_p, and Doppler width σ) are identical for each transition, as in [158], this shape function can then be inserted into equation 3.118 to obtain the total refractive index of the medium, and thus find the resulting cavity transmission function.

Having obtained the full expression for the cavity response including the contribution from the vapour, we note that the strength of light matter interaction is often discussed in terms of the optical depth, d, of the matter ensemble, where d is defined by the on resonance absorption of the medium: $I_{\text{out}} = I_{\text{in}} \exp(-2d)$.

Then for the bare atoms (i.e. neglecting any Doppler or pressure effects) perfectly initialised into the ground state i, from equation 3.119 on resonance we have that $s(0) = \mathrm{i}/\gamma$, so d can be expressed as [166]

$$d = \frac{1}{2}\frac{\omega_0 L}{c}\frac{N D_{if}^2}{\hbar\varepsilon_0}\frac{1}{\gamma}, \tag{3.125}$$

where N is the number density of caesium atoms interacting with the light field, in some length L. With this we can rewrite the susceptibility as

$$\chi = \frac{2d\gamma}{\omega_0 L/c} s_\lambda(\Delta). \tag{3.126}$$

Noting that far away from resonance, the Lorentzian character of the shape profile, $s(\Delta)$, is dominant [157], in the limit of large detuning the absorption and dispersion have the functional dependance

$$\lim_{\Delta \gg \gamma,\sigma} \left(\frac{\omega L}{c} n'\right) = \frac{\omega L}{c} - \frac{\omega}{\omega_0}\frac{d\gamma}{\Delta} \tag{3.127}$$

$$\lim_{\Delta \gg \gamma,\sigma} \left(\frac{\omega L}{c} n''\right) = \frac{\omega}{\omega_0}\frac{d\gamma^2}{\Delta^2}. \tag{3.128}$$

Noting that the dispersive part (n') of the complex refractive index decays slower ($1/\Delta$) away from resonance, this means that measurements sensitive to the dispersion yield information about the atomic population over a wider frequency range than measuring only the absorption (which scales as $1/\Delta^2$), as exploited in reference [135]. This affects the frequencies of the cavity resonances and

therefore provides both a means to measure the atomic populations and linewidths as well as to control the position of the full cavity transmission frequencies.

With this description, then, the intra cavity field and transmission can be obtained from 3.112 (via equation 3.117), as a function of physical parameters: the temperature of the vapour (broadening), the caesium reservoir temperature (number density), population distribution between the ground states (optical pumping efficiency) and the cavity reflectivities.

3.3.3 Dispersion Model Applied to the Cavity Memory

We now apply the general model developed in the previous section to our particular experimental arrangement for the cavity enhanced memory, in which we have a small vapour cell inside a ring cavity shown in figure 3.3a.

Note that the ring cavity configuration doubles the effective length of caesium compared to a single pass in free space. Another motivation for using a ring cavity for light matter interaction experiments using warm vapour is to avoid the formation of a standing wave as would form in a linear cavity where the intensity modulation over length scales of a wavelength would otherwise limit the coherence lifetime due to diffusion of the atoms out of the beam area, where at room temperature the atoms have a typical thermal velocity $\sqrt{k_B T/m} \sim 100$ ms^{-1}.

A schematic of the energy level structure of the D_2 line of caesium is shown in figure 3.3b, where the ensemble of atoms are optically pumped into one of the hyperfine groundstate levels. In the present analysis we consider the case of zero magnetic field for simplicity, and so we neglect consideration of the hyperfine sublevels. Linear polarisation is used for both the optical pump and signal light. To account for the effect of imperfect pumping of the atoms, however, when evaluating the susceptability we therefore sum the contributions of each transition to the overall susceptibility with their appropriate detunings and weighted according to their populations. The population distribution is defined in terms of the atomic polarisation, $w = p_1 - p_3$ where p_i are the population fractions of the hyperfine ground state levels i, and $-1 \leqslant w \leqslant 1$ quantifies the optical pumping efficiency.

For the purposes of the memory experiment the cavity is configured with a free spectral range of $\Delta\omega \sim 7$ GHz in the absence of the vapour, with an acceptance bandwidth $\delta\omega \sim 1$ GHz. Where we aim to meet the $M = 2$ resonance condition with an acceptance bandwidth compatible with our signal.

With this picture, equation 3.112 is modified slightly to become:

$$\frac{E_{\text{trans}}}{E_{\text{in}}} = \frac{t_{\text{in}} t_{\text{out}} \mathrm{e}^{\mathrm{i}\delta}}{1-\zeta}, \tag{3.129}$$

where now δ is the component of the phase acquired between the input and output mirrors where we assume no loss. From this, we can then extract expressions for the free spectral range (FSR, $\Delta\omega$), bandwidth ($\delta\omega$), and finesse (F).

Having chosen the FSR ($\Delta\omega = 2\pi c/L$) in order to meet the specific multi-resonances conditions above, appropriate reflectivities can be chosen such that the cavity bandwidth matches that of our input signal, with a FWHM of 1.2 GHz, via F

$$F = \frac{\Delta\omega}{\delta\omega} = \frac{\pi}{2\arcsin\left(\frac{1-\zeta}{2\sqrt{\zeta}}\right)}. \tag{3.130}$$

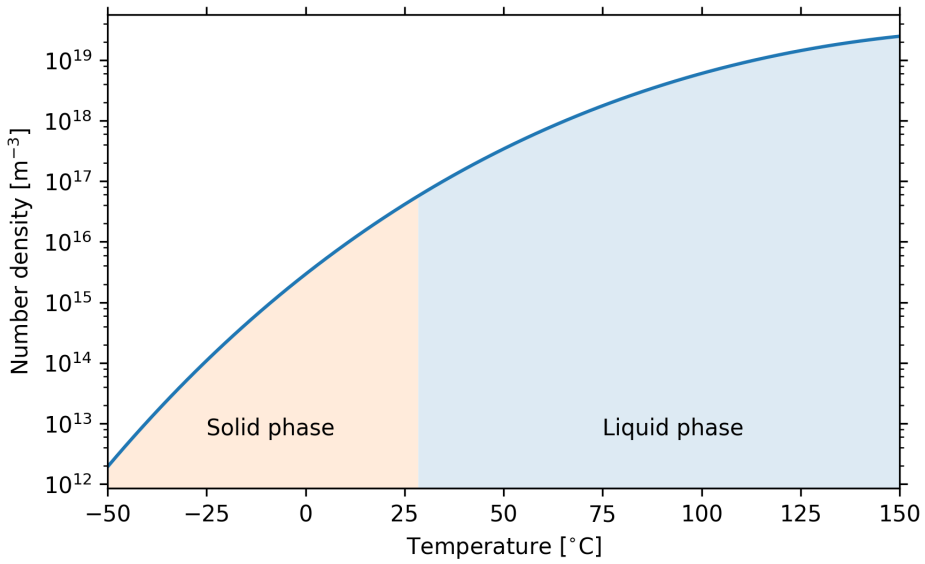


Figure 3.6: Estimated number density of caesium atoms from the vapour pressure as a function of reservoir temperature, evaluated from [143]

As before, separating the absorptive and dispersive parts, writing $\zeta = |\zeta| e^{i\phi}$, then gives

$$|\zeta(\omega)| = \zeta_0 \exp\left[-\frac{\omega}{c} n''(\omega) L_{\mathrm{Cs}}\right] \tag{3.131}$$

$$\phi(\omega) = \frac{\omega}{c}\left[n'(\omega) L_{\mathrm{Cs}} + L_{\mathrm{out}}\right], \tag{3.132}$$

where L_{Cs} is the total distance of caesium travelled through in one round trip, L_{out} is the additional path length outside of the vapour cell, and n is the complex refractive index of the caesium. Note that we have assumed that the contribution to ζ from the mirrors is purely real and is included in ζ_0, along with any other intracavity losses.

For the vapour cell used in this experiment, estimates for the pressure induced broadening and shift of the caesium transitions due to the buffer gas, at a given temperature was made using the results presented in reference [160, Tables II & III]. Here the self contribution to pressure broadening from caesium is negligible due to the low densities in our operating regime $\sim 10^{12}\,\mathrm{cm}^{-3}$ [167]. The total number density of caesium atoms, plotted in figure 3.6, in the vapour phase is estimated from the temperature of the caesium reservoir at the coldest point in the cell, using the expression for the vapour pressure, $P_{\mathrm{v}}(T)$, given in equation 1 of [143], via the ideal gas law $n_{\mathrm{v}}(T) = P_{\mathrm{v}}/k_{\mathrm{B}}T$. The reflectivities, transmissions and losses for the cavity optics were also characterised experimentally.

The frequency of a probe laser was scanned and the transmission through the cavity in the vicinity of the D_2 transitions measured. Low probe powers ($\sim \mu$W) are used so as not to affect the population distribution during the measurement.[14] Comparison with this model allows us to obtain an estimate of the physical parameters: critically the optical depth and the pumping efficiency which are key figures of merit for the memory interaction in [153].

[14] c.f. a saturation intensity of $\sim 2\,\mathrm{mW\,cm}^{-2}$ [143].

3.3.4 Numerics

The resulting function to which we fit the data is given in full by equations 3.133-3.134 below. Together, these define the cost function optimised when fitting the data.

$$\frac{I_{\text{out}}(\omega,\{T_k\})}{I_{\text{In}}} \approx A + B g_{\sigma_{\text{p}}}(\omega) * \left| \mathrm{e}^{\left[\frac{\mathrm{i}\omega L_{\text{out}}}{c}+\mathrm{i}\phi\right]} \left(1 + |\zeta_0|\, \mathrm{e}^{\left[\frac{\mathrm{i}\omega L_{\text{Cs}}}{c}\left(1+\frac{1}{2}\sum_{ij}\chi_{ij}(\omega,\{T_k\})\right)+\frac{\mathrm{i}\omega L_{\text{out}}}{c}+\mathrm{i}\phi\right]}\right)^{-1} \right|^2, \tag{3.133}$$

where

$$\begin{aligned} \chi_{ij}(\omega,\{T_k\}) &= c_{ij}^2 \frac{p_i n_{\text{v}}(T_{\text{res}}) D_{ij}^2}{\hbar\epsilon_0} s\left[\omega - \omega_{ij} - \delta_{\text{coll}}(T_{\text{P}}),\{T_k\}\right] \\ s(\Delta,\{T_k\}) &= \left[f_{\gamma+\gamma_{\text{coll}}(T_{\text{P}})} * g_{\sigma_{\text{D}}(T_{\text{D}})}\right](\Delta). \end{aligned} \tag{3.134}$$

The coefficients A and B are experimental parameters determined by the background and overall scaling, but the remaining constants can be used to extract the temperature via the Doppler width $\sigma_{\text{D}}(T_{\text{D}})$, the pressure broadening $\gamma_{\text{coll}}(T_{\text{P}})$ and shift $\delta_{\text{coll}}(T_{\text{P}})$, and the number $n_{\text{v}}(T_{\text{res}})$ of atoms, and populations p_i of the hyperfine ground state levels i.

In the preceeding expressions, we have distinguished three effective temperatures $\{T_k\}$, where the indices $k \in \{\text{'res'},\text{'P'},\text{'D'}\}$ refer to the temperatures corresponding to the caesium reservoir which determines the number density, pressure broadening and shift, and Doppler broadening. We permit these to vary independently as they are not necessarily equal [158] due to the non-uniform temperature within the cell and to account for the limited accuracy of the formulae for number density and pressure width. Additionally, we include a correction phase $\phi \in [-\pi,\pi]$ to the external cavity length L_{out}, where ϕ is determined by the setpoint of the cavity lock.

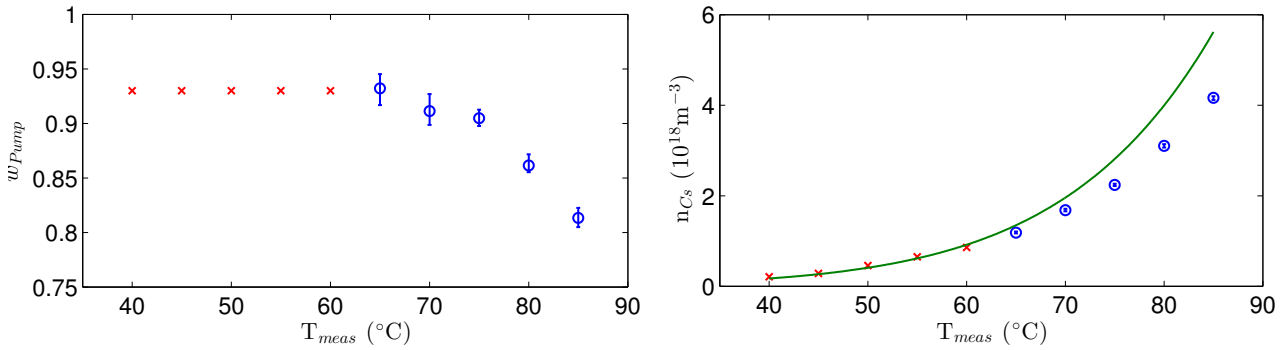
The functions $g_\sigma(\omega)$ and $f_\gamma(\omega)$ are defined as in equations 3.122 and 3.119. The convolution in equation 3.134 combines the homogeneous and inhomogeneous broadening terms, and in 3.133 accounts for the bandwidth of the probe laser. We note that as an alternative to performing the convolution of f and g explicitly, there exist algorithms for evaluating the Faddeeva function, see [164]. This method was also implemented and compared with the output from performing the convolution; however, both methods performed comparably, with the numerics for the convolution being slightly faster in our implementation.

Due to the number of parameters describing the system in equations 3.133-3.134, this problem is a difficult one to search numerically as the natural periodicity of the function results in many local minima, and each evaluation of the profile is itself intensive due to the two nested convolutions or the calls to the Faddeeva package. In order to make the problem more tractable two "one off" calibration measurements were performed aiming either to constrain variables or to provide favourable starting estimates for the fitting routine.

Firstly, in order to relate the measured surface temperature of the cell to the effective temperatures of the vapour, the reservoir temperature and buffer gas effects were estimated at a range of temperatures by measuring the absorption spectrum of the caesium in the absence of the cavity. Fitting the absorption spectrum is less numerically intensive than for the whole cavity spectrum and so we readily obtain values for the reservoir temperature and effective temperature of the motional

broadening at a given heater wire current. As indicated previously, treating these independently allows for differential temperature across the cell and uncertainties in the expressions used to define the widths and number density. We note that for the temperature range investigated, it was observed allowing the effective temperatures (T_{res}, T_{P}, T_{D}) to vary independently yields a better fit than defining a single temperature parameter. At lower temperatures it was observed that the obtained effective motional temperature (Doppler and pressure) is consistently greater than the reservoir temperature ($T_{\mathrm{res}} < (T_{\mathrm{P}}, T_{\mathrm{D}})$), whilst at higher temperatures both approaches gave comparable results ($T_{\mathrm{res}} \approx T_{\mathrm{P}} \approx T_{\mathrm{D}}$). The crossover between these regimes occurs at $\sim 60°\mathrm{C}$.

Secondly, with the cavity constructed around the cell, the roundtrip distance was constrained by measuring the cavity response far from the atomic resonance where dispersion is negligible. This constrains the cavity length ($L_{\mathrm{out}} + L_{\mathrm{Cs}}$) to within a wavelength, and fixes[15] the intracavity loss ζ_0. In the subsequent measurements, the total cavity length was then fixed by the correction phase ϕ such that the cavity is locked to maximise transmission at a detuning of $+6\,\mathrm{GHz}$.



(a) Pumping efficiency

(b) Caesium number density

Figure 3.7: Extracted parameters as a function of measured temperature. Blue circles indicate values obtained using the automated fitting routine, and red crosses those obtained using manual fits.3.7a optical pumping efficiency.3.7b caesium number density: the green line gives the number density expected if the caesium reservoir were held at the temperature measured at the cell surface. These appear in [154].

Based on the above, one approach for extracting parameters was to constrain the pressure and Doppler temperatures to equal that of the reservoir so as to reduce the size of the parameter space to be searched. The pumping efficiency and reservoir temperature, which together determine the optical depth, extracted with this method are plotted in figure 3.7. As observed in the measurement of the absorption spectrum, at higher temperatures ($\gtrsim 60°\mathrm{C}$) the assumption of a single temperature seems to reliably yield good fits, whereas at lower temperatures the fitting method became less robust. A typical fit obtained using this method is shown in figure 3.8b for measurements taken at the maximum operating temperature. The Peltier element was driven at a constant power and therefore the temperature differential across the cell varies with the current driving the heater wire temperature which was used to set the reservoir temperature. This may explain why the equal temperature approximation, required for the numerics to converge, breaks down at lower temperatures, and in this regime we rely upon manual searches.

Manual searches were performed by iterating over the Doppler and pressure widths and the reservoir temperature perturbed over some physically reasonable range from those obtained in fitting the absorption spectra at the same heater wire current, allowing for variation in environmental condi-

[15]The optical response of the optics are approximately constant over the relevant range of wavelengths.

Parameter	Manual	Automated
T_{Res} / °C	80.8	80.6
γ_{P} / MHz	41	65
σ_{D} / MHz	160	175
η	0.80	0.81

Table 3.3: Parameters extracted from the two fitting routines for data obtained at a measured temperature of 85°C at the cell surface.

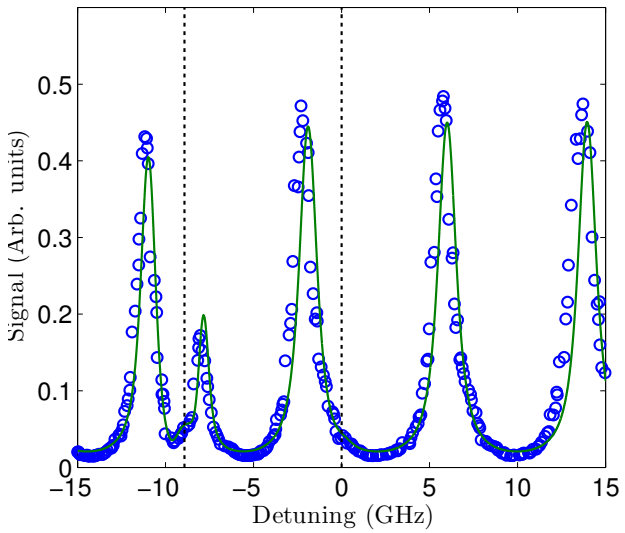
tions. The cost function was then optimised over the remaining unconstrained parameters, namely the pumping efficiency (w), background intensity on the detector (A) and the overall scaling of the photodiode measurements (B). For this approach we generate the modelled cavity response drawing parameters from a 6 dimensional array containing the perturbed parameters (T_{res}, σ_{D} and γ_{P}), and a range of estimates for A, B and w. From this, we then identify the set of parameters within the array yielding a minimum in residuals. Typical fits obtained using this method shown in figure 3.8a for measurements taken at the minimum operating temperature. We note that at a measured temperature of 85°C at the surface of the cell, we obtain the approximate values given in table 3.3, which are comparable to those obtained using the automated fits with a single temperature parameter.

To benchmark the performance of the fitting routine, we investigate the pumping efficiency and number density as functions of temperature, both of which are well studied [143, 168]. The pumping efficiency is seen to roll off at ~ 70°C, at which point increasing the temperature, and hence density of Cs, leads to light trapping impeding the ability to optically pump efficiently [168]; where, as the density of atoms is increased, the probability of a scattered photon exciting another atom increases (we return to this in § 4.3). The extracted reservoir temperature is in good agreement with the recorded temperature at the surface of the cell. Observation of this known behaviour confirms our numerical modelling. Error bars are generated for the automated fitting routine using the Monte-Carlo method over 100 trials. The limited pumping efficiency seen at temperatures below 60°C , as opposed to increasing to unity as temperature is decreased, observed in these data likely arises from a combination of imperfect mode overlap and limited pumping light entering the cavity.

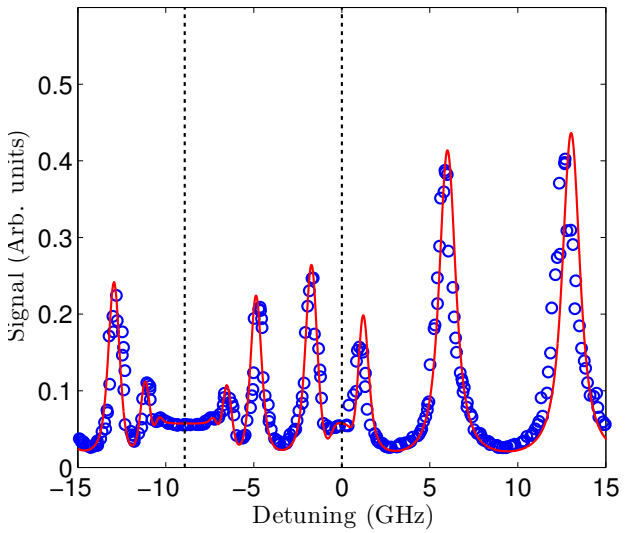
For the purposes of the cavity enhanced Raman memory, as in other experiments of atomic-like systems in cavities, control of the optical depth, pumping efficiency and cavity resonance frequencies are critical. In this section we have demonstrated this control, and investigated the variation of these key parameters as a function of temperature with caesium vapour in a cavity. These scans were recorded at 5 K intervals, and are shown in figure 3.9. The response of the cavity, falling pumping efficiency and increasing number density of caesium with temperature seen in these data are captured well by the model.

3.4 Implementation

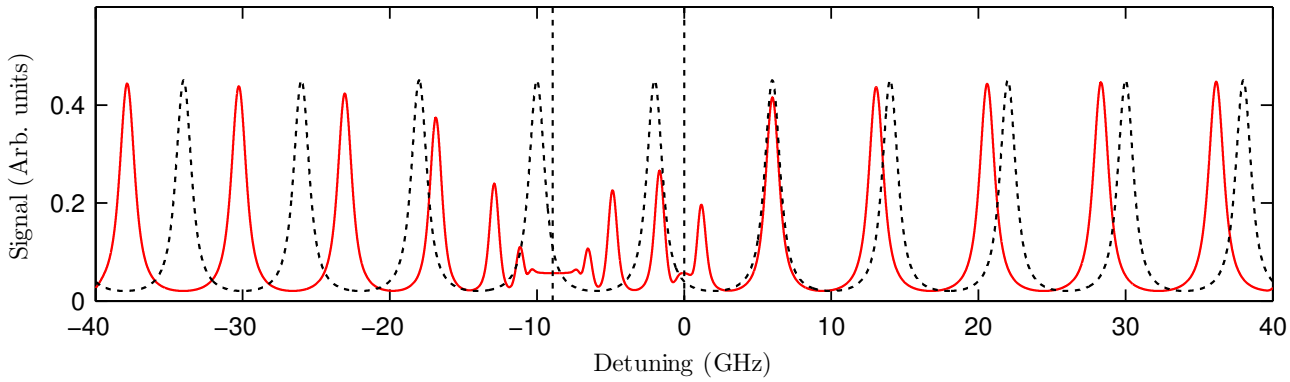
We now discuss the first demonstration of a Raman memory inside a cavity, published in [153]. With the previous ingredients we are able to simultaneously reach the required resonance conditions for the memory experiment, and from measurement of the optical response of the cavity extract the experimental parameters needed for our theoretical description of the memory experiment.



(a) Manual fit at 40° C



(b) Automatic fit at 85° C



(c) Comparison between response for filled and empty cavity.

Figure 3.8: Example of (a) manual search result at 40°C and (b) automatic fit at 85°C. The dashed lines indicate the frequencies corresponding to the centre of the two hyperfine resonances, $F = 4$ at -9.2 GHz and $F = 3$ at 0 detuning. Circles indicate experimental data, and the lines the theoretical response for manually fitted paramaters. (c) Comparison of modelled cavity response from automatic fit at 85° C - (as in (b), red) with cavity response with no atoms present (black, dashed). (Figures from [154])

3.4.1 Experimental Details

The overall scheme for the experiment is summarised in figure 3.10. The master laser for the experiment is a mode-locked Ti:Sa oscillator (Spectra Physics Tsunami) with an intracavity Gires-Tournois interferometer (GTI), emitting an 80 MHz pulse train of 360 ps, 1.2 GHz bandwidth pulses at 852 nm with an average power of 1.2 W. The Ti:Sa frequency is referenced to the pumping diode laser, and detuned at +15.2 GHz from the blue transition. The relative frequency of the two lasers are monitored on a scanning Fabry-Pérot (SFP) interferometer (Thorlabs) with an FSR of 10 GHz. The centre wavelength of the Ti:Sa was locked by actively feeding back on the GTI voltage based upon the SFP signal via LabView software originally written by a former student [169].

The pulse train is passed through a Pockels cell (Quantum Technology QC), with an extinction of 1000:1, in order to "pick" a series of 9 successive pulses in order to prevent pumping by the control field before storage, with an overall experimental repetition rate of 8 kHz.

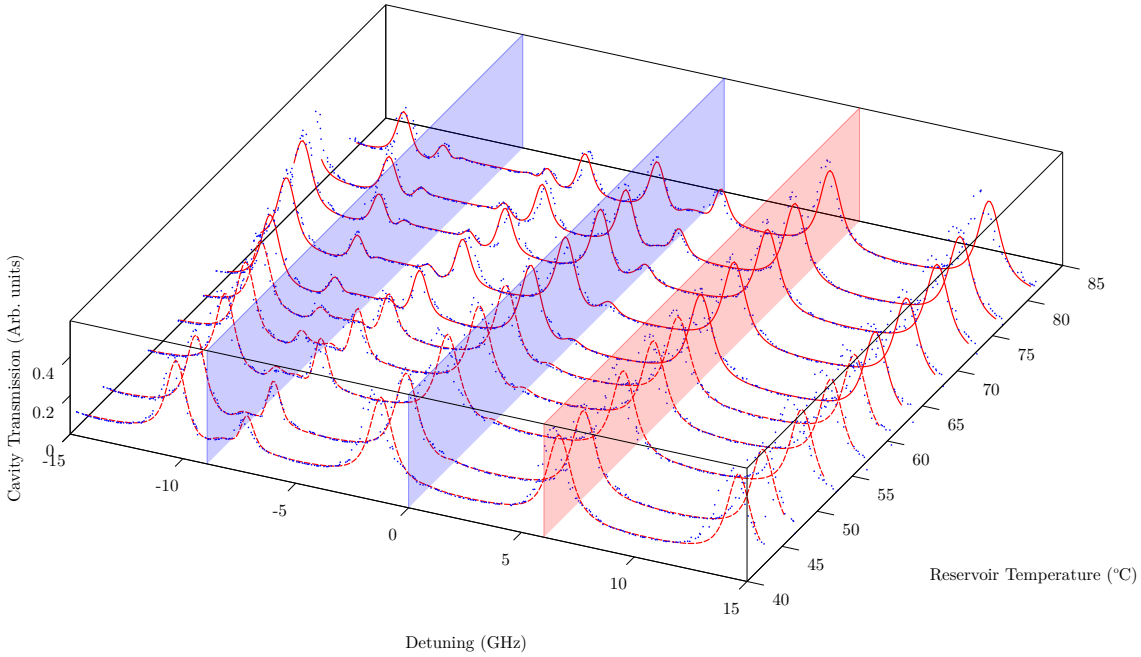


Figure 3.9: Cavity transmission, locked to resonance at 6 GHz blue detuned from the $F = 3$ resonance. Showing our model, from fitted parameters obtained from measurements at 5°C intervals. Blue surfaces indicate the frequencies of the two D_2 resonances, and the red surface the position of the cavity lock. Solid red lines indicate the model obtained using a fitting routine, red dashed lines show the manually obtained models (from [154]).

3.4.1(a) Signal Preparation

A polarising beam splitter (PBS) is used to pick off light to pass through an EOM generating 9.2 GHz sidebands on the first (write) pulse in order to produce the coherent state signal, which is filtered from the carrier by passing through a 36.8 GHz FSR etalon. The signal and control are then spatially and temporally overlapped, combined on a PBS and coupled into the same single mode fibre (SMF) and sent to the cavity. This set up is sketched in figure 3.10.

3.4.1(b) Optical Pumping

To optically pump the atomic ensemble, a CW ECDL (Toptica DL Pro) was locked to the crossover resonance $|^2S_{1/2}; F = 3\rangle \longrightarrow |^2P_{3/2}; F = 2\rangle$, $|^2P_{3/2}; F = 3\rangle$, and was sent through the cell, counter propagating along the control beam path. To ensure spatial overlap of the pumping region with the signal and control mode, the pump light was retro–coupled into the signal input SMF.

The linewidth (< 1MHz) of the pumping laser means that only atoms in a small velocity class are efficiently pumped: we therefore rely upon both (i) frequent collisions[16] in order to mix velocity classes and address all atoms in the ensemble; and (ii) the power broadening of the atomic transition. In order to facilitate memory lifetime measurements, discussed later, the pumping laser is passed through an AOM, and the first order diffracted beam path coupled to the experiment. This enables the pumping to be switched off during the memory protocol which avoids pumping out of the storage state which otherwise limits both the achievable memory efficiencies and lifetimes [166]. We note that when the AOM is included, the diffraction efficiency ($< 70\%$) limits the available power for optical pumping. It was noted by Michelberger *et al.* in the supplementary material for [118], that whilst

[16] Note that collisions with the neon buffer gas are spin preserving [166].

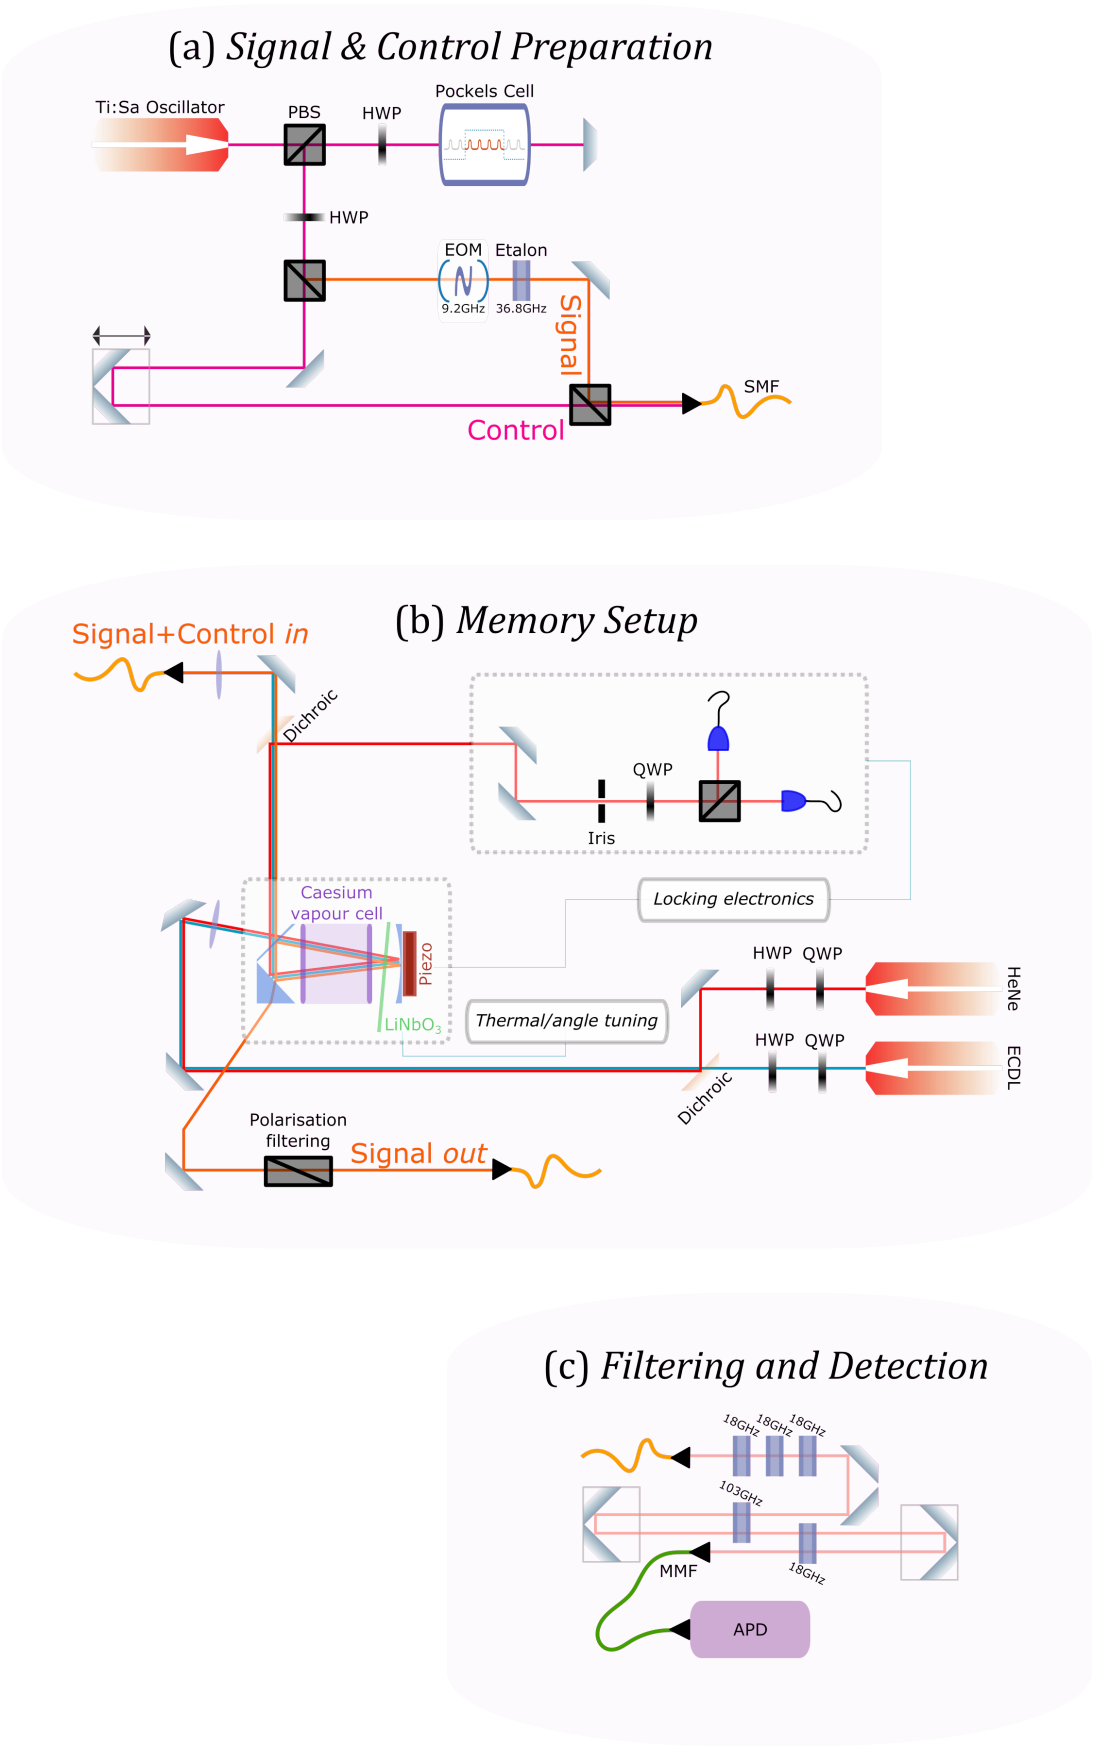


Figure 3.10: Experimental schematic for the cavity memory set up: (a) the preparation of coherent state signal and control pulse, derived from Ti:Sa oscillator, (b) the cavity memory experiment, including control system inputs, (c) the detection stage. (Note that a similar set-up is used for the results in section 7.3.2.)

the presence of the optical pumping during the memory protocol slightly reduced the overall memory efficiency it had no affect on the $g^{(2)}$ of the retrieved signal, nor the $g^{(2)}$ of the noise in both time bins, and so for measurements at fixed and short storage time (12.5 ns), the AOM can be removed, and the atomic polarisation can reach steady state, limited by light-trapping and depolarising collisions (cell wall, or Cs-Cs which occur on much longer timescales).

3.4.1(c) Filtering and Detection

In order to isolate the signal from other fields, the output of the cavity is filtered in polarisation, frequency and space; passing first through a Glan-Taylor polariser and a polarising beam displacer. The signal polarisation is then coupled into SMF leading to a series of air spaced Fabry-Perot etalons: three single passed etalons with an FSR of 18.2 GHz, then one double passed with FSR of 103 GHz, and final double passed with 18.2 GHz, giving an effective total of seven etalons.

The filtered signal is then detected with a single photon counting module (Perkin Elmer), and the arrival times logged with a time–to–digital converter (qutools quTAU), with a bin resolution of 81 ps. The overall transmission in the detection stage (from output of the fibre to the APD input in figure (c) of 3.10) is 1%.

3.4.2 Memory Results

The procedure for running the memory experiments involves first aligning the cavity whilst scanning the length attempting to maximise transmission and visibility and the mode quality with the laser locked far from resonance. The pumping alignment is optimised whilst still scanning the cavity length by maximising the visibility of the laser locked to the signal frequency. Finally the cavity length is locked to maximise transmission of the signal frequency, and by scanning the laser frequency the triple resonance condition is found by adjusting the cavity temperatures.

In the following section, we discuss the first series of measurements taken once *one* such solution for the triple resonance had been found.

3.4.3 Storage of Coherent States

3.4.3(a) Memory Efficiency Calculation

The efficiency of a quantum memory is defined as the probability of retrieving a photon given that a photon was *attempted* to be stored in the memory. For coherent states, this is simply expressed in terms of the average photon number *at the memory* in the input and retrieved pulses. For clarity, we define the notation n_i^L to denote the average number of photons *in the signal mode* in the time bin specified by the pulse time bin where $i = 0, 1, ..., 8$ (i.e. $i = 0$ corresponds to the *write* time bin, $i = 1$ to the first[17] readout time bin), and the combination of fields applied $L \in \{''\text{CSD}'', ''\text{SD}'', ''\text{CD}''\}$ in terms of Control (C), Signal (S) and optical pumping Diode (D); e.g. n_0^{CSD} corresponds to the number of photons in the signal mode arriving at the detectors in the first time bin when all three fields are applied.

[17] Since, in principle any signal that is not read out in the first readout bin, may be retrieved in subsequent bins, c.f. [170].

Assuming a noiseless memory, we would then have:

$$\begin{aligned}\eta_{\mathrm{m}} &= \frac{1}{n_0^{\mathrm{SD}}}\sum_{i=1}^{8} n_i^{\mathrm{CSD}} \\ &\approx \frac{n_1^{\mathrm{CSD}}}{n_0^{\mathrm{SD}}},\end{aligned} \tag{3.135}$$

where for now, we neglect the contribution from subsequent read out ($i > 1$). This can be decomposed into *write* and *read* efficiencies:

$$\eta_{\mathrm{write}} = \frac{n_0^{\mathrm{SD}} - n_0^{\mathrm{CSD}}}{n_0^{\mathrm{SD}}}, \tag{3.136}$$

$$\eta_{\mathrm{read}} = \frac{n_1^{\mathrm{CSD}}}{n_0^{\mathrm{SD}} - n_0^{\mathrm{CSD}}}, \tag{3.137}$$

$$\eta_{\mathrm{m}} = \eta_{\mathrm{write}}\,\eta_{\mathrm{read}}. \tag{3.138}$$

In practice, we must account for both (i) noise in the signal mode (from FWM); and (ii) spurious counts from leakage of other modes.

Figure 3.11 shows a typical run of the memory experiment, with a memory efficiency of ∼8%, for a coherent state with an average of ∼20 photons. The respective counts in each bin, for each combination of fields is found by integrating under the appropriate curve within the shaded regions.

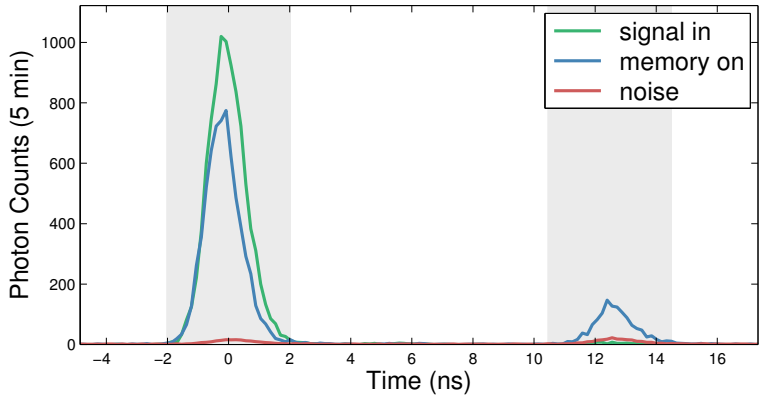


Figure 3.11: Typical plot of time tagger data for memory experiments. Here the legends correspond to "signal in": SD, "memory on": CSD, "noise": CD.

In order to make a measurement of the n_i^L, the output of the cavity is sent through a polarisation and frequency filtering stage to an APD.[18] The output from the APD was logged with a time-to-digital time-tagger, and logged using a custom software interface written by A.E.. In order to relate the number of *signal* photons n_i^L to the total number of "clicks" at the APD, which we denote N_i^L, one must account for imperfect filtering, dead time, dark counts.

Accounting for imperfect filtering, we have photons arriving at the detectors in the Signal/Stokes (S), control (C) and anti-Stokes (A) modes and some contribution from fluorescence (Fl.), each of which experience a different transmission efficiency between the memory and the detector. We denote the number of photons at the detectors in each mode $M \in \{\mathrm{S}, \mathrm{C}, \mathrm{A}\}$. Given perfect filtering we would

[18] This was eventually intended to be pair of APDs, one at each output port of a 50:50 beam splitter which would (a) reduce the effective deadtime of the detection and (b) enable later measurements of photon autocorrelation statistics. However, for the first measurement a single APD is sufficient.

have that $N_i^L(M \neq \mathrm{S}) = 0$, but in general the number of clicks is instead

$$\begin{aligned} N_i^L &= N_i^L(\mathrm{Fl.}) + N_i^L(\mathrm{S}) + N_i^L(\mathrm{C}) + N_i^L(\mathrm{A}) \\ &\approx N_i^L(\mathrm{S}) + N_i^L(\mathrm{C}), \end{aligned} \tag{3.139}$$

where we have assumed that with sufficient filtering, the only significant leakage contribution at the detectors should only be from the strong control field.[19] So, to characterise the memory we can consider only FWM noise in the signal mode, the true signal, and control field leakage. Additionally, the EOM which generates the signal has a finite switching time which can give a nonzero input signal in the read out time bins.

In order to analyse the memory performance, we must relate the number of clicks at the detector to the number of signal photons at the memory output. With the levels of suppression required in the filtering stage in order to conduct single photon level experiments, it is not feasible to directly characterise the filtering efficiency of unwanted modes directly. Nonetheless, with this formalism, the number of clicks at the detector due to the true signal when the memory is on (control field applied) is given by:

$$\begin{aligned} N_i^{\mathrm{CSD}} - N_i^{\mathrm{CD}} &= \left[N_i^{\mathrm{CSD}}(\mathrm{S}) + N_i^{\mathrm{CSD}}(\mathrm{C})\right] - \left[N_i^{\mathrm{CD}}(\mathrm{S}) + N_i^{\mathrm{CD}}(\mathrm{C})\right] \\ &= N_i^{\mathrm{CSD}}(\mathrm{S}) - N_i^{\mathrm{CD}}(\mathrm{S}), \end{aligned} \tag{3.140}$$

where we have assumed that the control leakage is the same in both cases: $N_i^{\mathrm{CSD}}(\mathrm{C}) = N_i^{\mathrm{CD}}(\mathrm{C})$.

The memory efficiency can then be evaluated as

$$\eta_{\mathrm{m}} \approx \frac{N_1^{\mathrm{CSD}} - N_1^{\mathrm{SD}} - N_1^{\mathrm{CD}}}{N_0^{\mathrm{SD}} - N_0^{\mathrm{CD}}}. \tag{3.141}$$

3.4.3(b) Memory Measurements

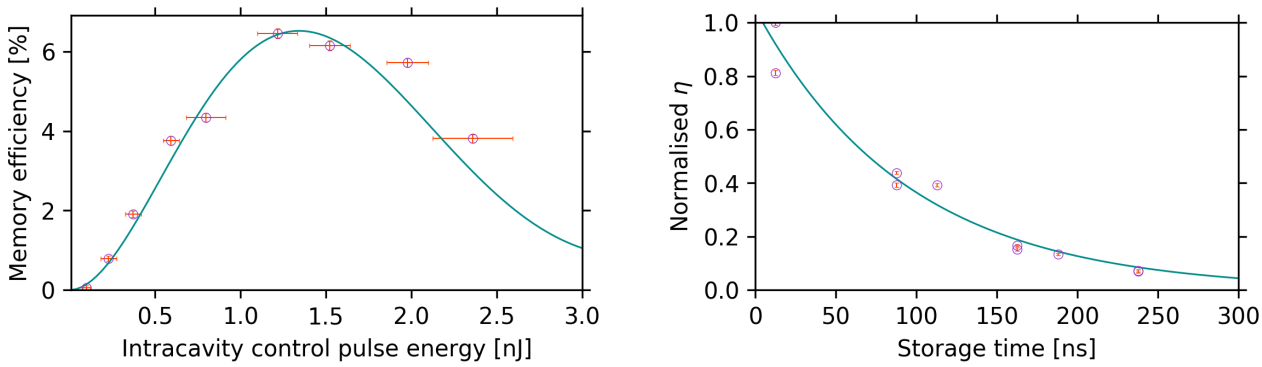
As discussed in section 3.3.3, by controlling the temperatures in the vapour cell, optical pumping efficiency, cavity birefringence, and the overall cavity length, the system can be engineered to meet the "triple resonance" condition required for memory operation. Once this was achieved, with the cavity length and temperature locked to maximise transmission of the signal mode, the memory efficiency was then further optimised upon varying the focussing into the cavity, mode overlap between signal and control (in time, space, and frequency). Once aligned, we measure the cavity visibility $V = (I_{\max} - I_{\min})/(I_{\max} + I_{\min})$ for the three relevant fields[20], where as discussed previously, the visibility is a critical factor in determining the suppression attainable by the cavity.

Field	**Visibility** [%]
Signal	86 ± 2
Control	71 ± 5
Anti–Stokes	86 ± 5

Having since developed the more rigorous description of the cavity response, for the measurements in figures 3.11 and 3.12 we estimate the cell temperature to be $\sim 85°\mathrm{C}$ with population inversion

[19] i.e. we neglect contributions from spontaneous Raman scattering and fluorescence, both of which are emitted into 4π, and so are highly suppressed by spatial and frequency filtering.

[20] Bearing in mind the polarisation and frequency dependence of the intracavity optics.



(a) Variation of memory efficiency with control power at preparation stage ($N_{\text{in}} = 0.7 \pm 0.1$).

(b) Memory efficiency with varying delay between read and write pulses ($N_{\text{in}} \sim 10$).

Figure 3.12: Results for memory efficiency with (a) varying control power, and (b) varying storage time. y error bars from Poissonian noise.

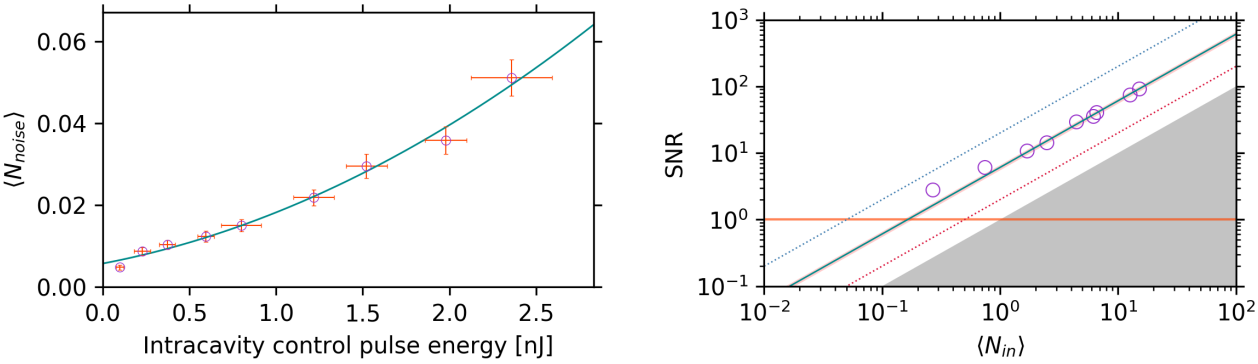
$w \sim 0.65$, giving an optical depth $d \approx 300$, and pressure broadened linewidth $\gamma \approx 250\,\text{MHz}$. The higher cell temperature than previous experiments (conducted at 70°C) increases the single pass optical depth of the ensemble, d, and hence the achievable memory efficiency. However, the combination of lower intracavity pumping power and higher cell temperature together limit our ability to effectively initialise the ensemble, in comparison with the state preparation efficiency of 95% (or, $w \sim 0.9$) in the previous, bulk experiment [118].

Before discussing the details of the series of experiments conducted, we first note the asymmetry in the peak shapes obtained from the time tagger data, e.g. figure 3.11; in addition to the dynamics due to the cavity, there is also some contribution to the group delay from the filtering stage, where the transmitted light is sent through an array of etalons in order to isolate the signal mode. This means that the extracted parameters, such as memory efficiencies and signal to noise ratio (SNR) depend on the integration window chosen.

The first series of measurements were therefore to investigate the memory efficiency for storage of weak coherent states. Due to the sensitivity of the system to varying the number density of caesium and its population distribution, the simplest measurements to begin to characterise the system is the overall memory efficiency function of control field power and storage time, shown in figure 3.12. Other insightful measurements such as variation with pumping efficiency and temperature are less trivial since at each stage a new solution to the triple resonance condition would have to be found.

Results for efficiency scanning the intracavity control field pulse energy, shown in figure 3.12a, show the characteristic memory behaviour as a function of Raman coupling. At low powers, we recall the the Raman coupling C, from equation 3.35, $C \sim |\Omega_{\text{max}}|$, which determines the maximum achievable memory efficiency (via the cooperativity parameter) hence we see that the memory efficiency initially increases with incident control field power. At elevated powers we begin to observe the contribution from the dynamic Stark shift, we have that the memory efficiency is expected to saturate at some control field intensity and then roll off with further increase, as observed in the data.

Measurements of the memory lifetime were conducted by evaluating the memory efficiency, whilst varying the delay between the write and read pulses. Fitting the achieved memory efficiency to an exponential decay yields a memory lifetime of $\tau = 95 \pm 7$ ns. This is comparable with the lifetime of $\tau = 150$ ns observed in the previous bulk experiment before magnetic shielding was included [166]. The discrepancy between the present measurement and the previous measurements could also have



(a) Noise counts scaling varying control field power.

(b) μ_1 parameter – extracted from the scaling of SNR with input photon number.

Figure 3.13: (a) Noise counts scaling with control pulse energy, with polynomial (quadratic) fit; (b) evaluation of the μ_1 parameter for the memory: dashed blue – optimal, dashed red – free space; error bars within the marker size. The intersection point with the horizontal line (where SNR = 1) yields the value of μ_1; where for a *practical* quantum memory $\mu_1 \ll 1$ are necessary. The shaded grey region indicates a the region yielding $\mu_1 > 1$.

a contribution from a worse extinction of the AOM as well as the smaller beam waist. Note, however, that once magnetic shielding was added to the bulk, a lifetime of $\tau = 1.5$ μs was then reached [107] and, based on their theoretical analysis of the dephasing mechanism due to stray magnetic fields (supplementary information [107]), the lifetime could have, in principle, been increased further with improved magnetic shielding. This is therefore an additional design feature which should be incorporated into subsequent generations of the cavity experiment.

We note that this first series of measurements was conducted at higher temperatures and lower pumping efficiencies than previous generations of the experiment. Inefficient initialisation of the ensemble may result in spontaneous Raman scattering into the signal mode; so, in future, operating at lower temperature to reduce light trapping should enable higher pumping efficiency [168]. Some investigation into improvement of the pumping efficiency is discussed in the next chapter.

Crucially, the results so far show that the noise retrieved from the cavity memory does not build up between subsequent applications of the control pulse: build up of the retrieved signal is a characteristic signature of FWM noise which was observed during experiments in the bulk caesium cell [146].

3.4.4 Noise Analysis

Measuring the scaling of noise counts in the read–out time bin as the control pulse energy, $\mathcal{E}$ is varied, shown in figure 3.13a, we may interpret three principle components: (1) constant background; (2) linear in $\mathcal{E}$; and (3) quadratic in $\mathcal{E}$.

The constant and linear terms were measured, and correspond to $\langle n^{\text{noise}} \rangle = 0.015 \pm 0.002$ photons per control pulse, at the control field energy corresponding to the maximal measured efficiency. This constitutes the noise floor for the memory operation.

To consider the impact of the cavity memory on the suppression of FWM, the contribution to noise counts from FWM are expected to scale quadratically with $\mathcal{E}$, i.e. the FWM process requires the scattering of two control pulse photons for the generation of a single noise photon in the signal mode. Analysing the scaling with control field power, and evaluating at the conditions giving maximal efficiency in figure 3.12a, we find a contribution of 0.006 ± 0.003 photons from four-wave mixing which

agrees favourably with our theoretical prediction of 0.005 – evaluated using equation 3.83, in the weak coupling limit and setting $N_{\mathrm{in},0} = 0$, which scales quadratically with the control field energy (since $|C_{\mathrm{s,a}}|^2 \propto W$).

Noise contributions linear in the control intensity arise from (i) control field leakage through the filtering stage; and (ii) spontaneous Raman scattering from un-pumped population in the storage state into the signal mode. Finally, the constant background arises from sources such as dark counts or fluorescence.

To characterise the noise level before moving towards photon counting measurements, we consider the μ_1 parameter as a metric to quantify the performance: where $\mu_1 = \langle n^{\mathrm{noise}} \rangle / \eta^{\mathrm{tot}}$ (see equation 1.2) [83], in other words the input photon number required to achieve a signal-to-noise ratio of unity.

In figure 3.13b, we show the scaling of the SNR with the number of input photons $\langle N_{\mathrm{in}} \rangle$ measured at the control field energy ~ 1.4 nJ corresponding to the maximal SNR measured in this experiment, from this we can then find the μ_1. These analyses were performed by P.M.L. and D.J.S., which show by fitting the data to $\mathrm{SNR} = \langle N_{\mathrm{in}} \rangle / \mu_1$ based upon the form of equation 3.85 for fixed control field parameters, we extract $\mu_1 = 0.17 \pm 0.02$, where the uncertainty was estimated using Monte Carlo simulation. We compare these data with the optimal case (blue, dashed) in which we predict $\mu_1^{\mathrm{ideal}} = 0.05$ (again from equation 3.85), and the free space memory (red, dashed) taking $\mu_1' \approx 0.5$ based on the measured SNR in our previous cavity free Raman memory implementation [118].

For a practical quantum memory, we must minimise μ_1; such that a sufficient signal-to-noise ratio can be obtained for incident photons with a finite heralding efficiency, for instance, when the average number of photons arriving at the memory is $\langle N_{\mathrm{in}} \rangle < 1$.

In the next chapter, we will build on these results and discuss optimisation stratagies based upon experimental and numerical investigations.

3.5 Summary

In this chapter I have followed three themes of investigation of the cavity enhanced Raman memory.

§3.2 Firstly, the development of a theoretical description of the dynamics, and deriving expressions for key figures of merit for the memory performance (see also [147]).

§3.3 Following that, I discussed a numerical model for the optical response of the atom filled cavity for the purposes of characterising the experimental system, and to enable the simulation of different configurations to explore how specific resonance criteria may be met (see [154]).

§3.4 Finally, I have presented the proof-of-principle results of the cavity Raman memory (see [153]).

CHAPTER 4

CAVITY MEMORY OPTIMISATION

In this chapter, I build on the published results of the previous chapter; and outline some of the work I have undertaken experimentally and numerically to help determine the future directions of the project. The diffusion measurements and analysis were all conducted by myself with some experimental assistance from K.T.K.; also for the numerical exploration of the consequences of different cavity designs, with some guidance & advice from J.N. and D.J.S. The investigations of memory efficiency with different buffer gases were conducted by S.E.T.

4.1 Introduction

Moving forward from the results of the preceding chapter, one notices that the objective of realising a practical optical quantum memory in this system relies upon optimisation based upon three key metrics (i) Efficiency; (ii) Lifetime; and (iii) Noise. It has been identified that the choice of buffer gas is critical for all of these aspects, particularly as the size of the cell and beam waist is reduced in scaling down the size of the memory device. Furthermore, it is clear that the ability to efficiently initialise the ensemble is also vital, which also relates to the choice of buffer gas.

The practice of introducing buffer gases in vapour cell atomic physics experiments is commonplace. The primary objective is to enhance the coherence times that can be achieved by slowing the diffusion of the atoms of interest, due to collisions with the buffer gas, hence, the foremost criterion for a candidate species is that collisions between it and the host atoms are spin preserving (i.e. that the buffer gas has zero spin), and so for instance, noble gases are a common choice. In addition to this, other factors that can be considered for a given buffer gas are their size and the pressure broadening/shift they induce in the host atoms. As has been noted in the previous chapter, the cavity response is sensitive to lorentzian pressure broadening far from resonance, particularly for the roundtrip phase accumulated, but also present in the linear absorption term – and hence contributes to the intracavity loss.

4.1.1 Memory Efficiency

Whilst characterising the cavity memory, analysis of the cavity spectrum determined that we achieved a population distribution of $\sim 80\%$ in the desired ground state, where the pumping efficiency achieved is in part limited by the ability to optically pump in situ. Furthermore, the limited pumping efficiency restricts the optical depth which can be achieved because of the inability to optically pump at high

temperatures arising from light trapping [171]. Another project which was conducted in parallel with the cavity memory experiment, led by S.E.T. was to improve the spin polarisations that can be achieved in the "bulk" (our terminology for the cavity free vapour cell) aiming to enhance the memory efficiencies achievable by reaching higher optical depths, and reduced Raman scattering from the remaining population in the undesired state [142], with a view that these results can be applied to future generations of the cavity memory.

Additionally, as mentioned in the previous chapter, noise counts after the memory may arise from spontaneous Raman scattering due to imperfect initialisation of the ensemble, and as such has previously limited the optical depths at which we could operate due to radiation trapping drastically impeding our ability to optically pump. The improvements in optical pumping demonstrated in this work should also lower the noise floor of our future experiment.

Some conclusions and consequences of this are discussed in section 4.3

4.1.2 Memory Lifetime

The limiting factors which determine the memory lifetime in the Raman scheme in warm vapour are (i) magnetic dephasing of the hyperfine levels due to stray magnetic fields; and (ii) diffusion of the caesium atoms into – and out of – the beam area,[1] in the time between initialisation of the atoms and the final read out. At longer timescales there is also an additional decoherence mechanism arising from Cs–Cs collisions which are not spin preserving. The overall relaxation rate is then a sum over the different relaxation channels

$$\begin{aligned}\frac{1}{\tau} &= \sum_j \frac{1}{\tau_{\mathrm{j}}} \\ &= \frac{1}{\tau_{\mathrm{mag.}}} + \frac{1}{\tau_{\mathrm{dif.}}} + \frac{1}{\tau_{\mathrm{col.}}},\end{aligned} \tag{4.1}$$

where the index j refers to magnetic dephasing ("mag.") diffusion ("dif.") and caesium-caesium collisions ("col."). The relaxation from caesium collisions is negligible due to the low densities (we estimate $\tau_{\mathrm{col.}} \sim 1\,\mathrm{ms}$), and that the 95 ns lifetime observed in the cavity-Raman memory demonstration of the previous chapter is a combination of diffusion and magnetic dephasing limited.

The introduction of magnetic shielding, previously demonstrated in the bulk, was shown to increase the achievable memory lifetime from $\sim 150\,\mathrm{ns}$ to $\sim 1.5\,\mu\mathrm{s}$ [166]. When moving to the cavity configuration, a smaller beam waist – and hence enhanced Rabi frequency, and increased coupling strength – can be achieved, however this therefore means that the contribution of atomic diffusion becomes much more critical. For instance, our demonstration of the cavity enhanced memory a lifetime of $95 \pm 5\,\mathrm{ns}$ was shown [153], by including magnetic shielding, it is anticipated that this can be increased.

Recently S.E.T. has undertaken an investigation of the limitations of the memory lifetime in the bulk, with three concentric cylinders of $\mu-$metal magnetic shielding. This showed a strong dependence of lifetime on the beam waist used [142] for waists $110 \rightarrow 165\,\mu\mathrm{m}$ giving lifetimes on the timescales of μs. This indicates that with such shielding, the lower limit to the magnetic-dephasing timescale is $\tau_{\mathrm{mag.}} > \mu\mathrm{s}$, and that in this case the diffusion rate of the caesium is a limiting factor for the memory lifetime in our target operating regime of achieving lifetimes exceeding $1\,\mu\mathrm{s}$ in the cavity.

[1] Recall, the thermal velocity of caesium atoms $\sim 100\ \mathrm{ms}^{-1}$. Also technically applicable to the decoherence of the spin-wave itself, although the wavelength of this is on a larger lengthscale.

We see that moving to the cavity memory, and with a view to scaling it down in the future, we must control the rate of diffusion. In section 4.2 I present an experimental study (and discussion of the implications) of the diffusion rate of caesium in different buffer gases at a range of pressures.

4.1.3 Noise Reduction

The principle of suppression of four-wave mixing noise in the Raman memory scheme by means of a cavity, outlined in the previous chapter, is intuitively straightforward. The essential figure of merit for how much suppression is possible with a given cavity configuration (for operation at a given temperature and detuning) is encapsulated in the finesse of the cavity: a higher finesse results in a greater suppression of the anti-Stokes mode with respect to the signal. For a given cavity length, then, this requires that the acceptance bandwidth of the cavity be reduced to increase the amount of suppression that can be achieved ($x \sim \mathcal{F} = \text{FSR}/\text{bandwidth}$). Conversely, for some chosen bandwidth, moving towards a smaller cavity (with a larger free spectral range) can also increase the finesse, but is technically more challenging to implement, and, at higher finesse, as explored later in section 4.4 the system efficiency becomes much more sensitive to intracavity loss. Although, in the future, it would be desireable to be able to implement a machined cavity with a length on the order of a few mm, which with careful engineering based on the following investigations, could provide a more scalable platform for the multiplexing of heralded single photons.

Based on the results of the previous chapter, and analysis of the scaling of the noise with control field power, we have demonstrated reduction of the contribution from four-wave mixing. With a view to improving upon this there are a number of approaches to increase the noise suppression further – in this chapter we seek to explore this parameter space, and the unavoidable trade off with other factors such as the bandwidth, lifetime and efficiency.

In order to aid this search, in section 4.4, I use the theoretical description presented in the previous chapter (section 3.2), and discuss some numerical exploration of the cavity enhanced Raman memory I performed, aiming to be a design tool to guide future work.

4.2 Buffer Gas (1): Atomic Diffusion Measurements

As argued in the opening of this chapter, 4.1, it is clear that in order to simultaneously optimise the efficiency, lifetime and the noise performance, that the choice of buffer gas species and the pressure at which it is introduced is critical. In this section, we seek to compare the diffusion properties with the buffer gases neon (as has been used historically by our group) with two alternatives: xenon, a heavier noble gas which gives a much lower diffusion rate, but also significantly increased pressure broadening, and nitrogen, which has comparable diffusion and broadening rates as neon, but is also a candidate for quenching which can be used to reduce light trapping and enable higher densities to be achieved with a higher degree of spin polarisation, which is discussed in the next section.

In the following discussion, we must bear in mind the trade off between the achievable diffusion-limited lifetime $\tau_{\text{dif.}}$ and the consequences of increased pressure broadening as more buffer gas is introduced.

The homogeneous broadening of the caesium D_2 line due to the introduction of different buffer gases can be parametrized in terms of the pressure broadening rate γ for a given buffer gas species,

Species	γ [MHz/Torr]
N_2	19.18 ± 0.06
Ne	9.81 ± 0.06
Xe	18.41 ± 0.07

Table 4.1: values for the pressure broadening coefficient at a temperature of 313 K for the three buffer gases presently considered, quoted from [160].

where the linewidth of the transition depends on pressure, P, expressed in [160] as:

$$\Delta\nu_{\mathrm{L}} = \gamma P + \gamma_{\mathrm{N}}, \tag{4.2}$$

where $\Delta\nu_{\mathrm{L}}$ is the homogeneously broadened width, γ_{N} the natural linewidth and γ the broadening rates, tabulated in table 4.1.

In order to investigate the optimal buffer gas species and pressure, the diffusion properties of caesium in the presence of various buffer gases (neon, xenon, nitrogen) were measured at room temperature. When the mean free path of the atoms is dominated by collisions with a buffer gas[2] the motion of the atoms may be described by a diffusion equation

$$\frac{\partial \varrho(\boldsymbol{r},t)}{\partial t} = D\nabla^2 \varrho(\boldsymbol{r},t), \tag{4.3}$$

where D is the diffusion constant, and $\varrho(\boldsymbol{r},t)$ describes the spatial distribution of the atoms, at position $\boldsymbol{r}$ and time t.

To obtain values for the diffusion coefficients, we follow the method described in reference [172], whereby the diffusion of the atoms is registered by the spatial and temporal variation in optical depth of the ensemble following optically pumping a small ($\sim 1\,\mathrm{mm}$) cylindrical region of the atomic vapour. After some programmable delay, the spatial absorption of a large ($\sim 1\,\mathrm{cm}$) counter propagating probe beam, on resonance with the atoms, is imaged. Atoms which have been optically pumped are transparent to the probe light, and thus give rise to a region of decreased optical depth and higher transmission of the probe light. By measuring the transverse spatial profile of the optical depth as a function of time delay between the pump and the probe we therefore can obtain information about the rate of diffusion of the caesium atoms.

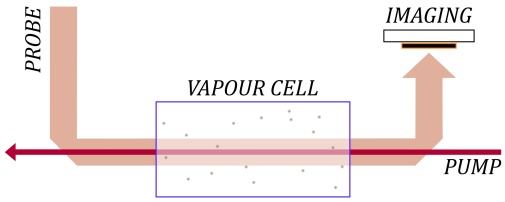


Figure 4.1: Pump probe configuration for diffusion measurements. The atomic ensemble is optically pumped for a fixed duration by a narrow collimated beam, and after a variable delay a wide counter-propagating probe beam is passed through the ensemble and imaged on a camera.

4.2.1 Method

The on–resonance absorption of light passing through the ensemble is related to the population distribution of the ensemble through the optical depth, d, via $I = I_0 \exp(-2d)$, where I_0 is the incident

[2] i.e. a length scale much smaller than the cell dimensions.

intensity profile. The optical depth is proportional to the number of atoms in a state resonant with the light field, over the length of the ensemble. Operating at room temperature, the number densities of caesium and thermal Doppler broadening of absorption features are small enough with respect to the ground state splitting such that we can neglect any linear absorption from atoms in the off–resonant ground state.

We therefore seek to obtain the transverse spatial and temporal dependence of the variation of d:

$$\begin{aligned}\Delta d(\boldsymbol{\rho}, t) &= \ln\left(\frac{I_{\text{pumped}}(\boldsymbol{\rho}, t)}{I_{\text{not pumped}}(\boldsymbol{\rho})}\right) \\ &= d_{\text{not pumped}}(\boldsymbol{\rho}) - d_{\text{pumped}}(\boldsymbol{\rho}, t),\end{aligned} \tag{4.4}$$

where $\boldsymbol{\rho}$ is the transverse coordinate and t the delay between switching the pump light off and applying/imaging the probe. This therefore directly yields the spatial profile of the population:

$$\Delta d \propto c_{\text{not pumped}}(\boldsymbol{\rho}) - c_{\text{pumped}}(\boldsymbol{\rho}, t), \tag{4.5}$$

where $c(\boldsymbol{\rho}, t)\,\mathrm{d}x\,\mathrm{d}y$ is the number of atoms at pixel $\boldsymbol{\rho}(x, y)$ per unit length in the atomic state resonant with the light, at some time t after the pump light has been applied.

For each cell, we therefore perform a pump-probe measurement measuring the spatial absorption profile of the light. In this case both pump and probe are resonant with the same transition - i.e. the optical pumping results in a local increase in transmission, this region then spreads out as a function of delay between the pump and probe. We record the transmission of the probe light (large diameter, low intensity) first through an un-pumped ensemble, and then for a series of delays after applying a pump pulse (small diameter, high intensity), see figures 4.2 and 4.3 for examples of typical data obtained.

4.2.2 Experimental Details

The pump and probe derived from the same diode laser (Toptica DL Pro) locked to the D_2 absorption line. The pump and probe are each switched on and off with an acousto-optic modulator (AOM) coupled into a single mode fibre (SMF) before the experiment. Both beams were collimated through the vapour cell (10 cm length, Precision Glassblowing) containing each buffer gas, for a series of pressures. The probe beam imaged onto a CCD camera (UniBrain). The 1 μs pulses permitted through the AOMs and data acquisition from the camera were triggered from a digital delay generator (Stanford Instruments DG645), controlled via MatLab. For each delay configuration, 100 images with an exposure of 1 μs were acquired. These are then averaged and processed.

A counter propagating configuration and orthogonal polarisations between the two pulses was used so as to seperate the intense pump and and weak probe light. Additionally, the short probe pulse is used so as to not dynamically affect populations whilst imaging.

4.2.2(a) Gaussian Approximation

To obtain an initial estimate for the diffusion coefficient, and give an intuitive picture of the physics, we assume an ideal Gaussian pump profile which is then convolved with thermal motion of the atoms. This results in a Gaussian profile in the radial coordinate, $\boldsymbol{\rho}$, with width σ_ρ

$$\sigma_\rho(t) = \sqrt{\sigma_\rho(t=0)^2 + 4Dt}. \tag{4.6}$$

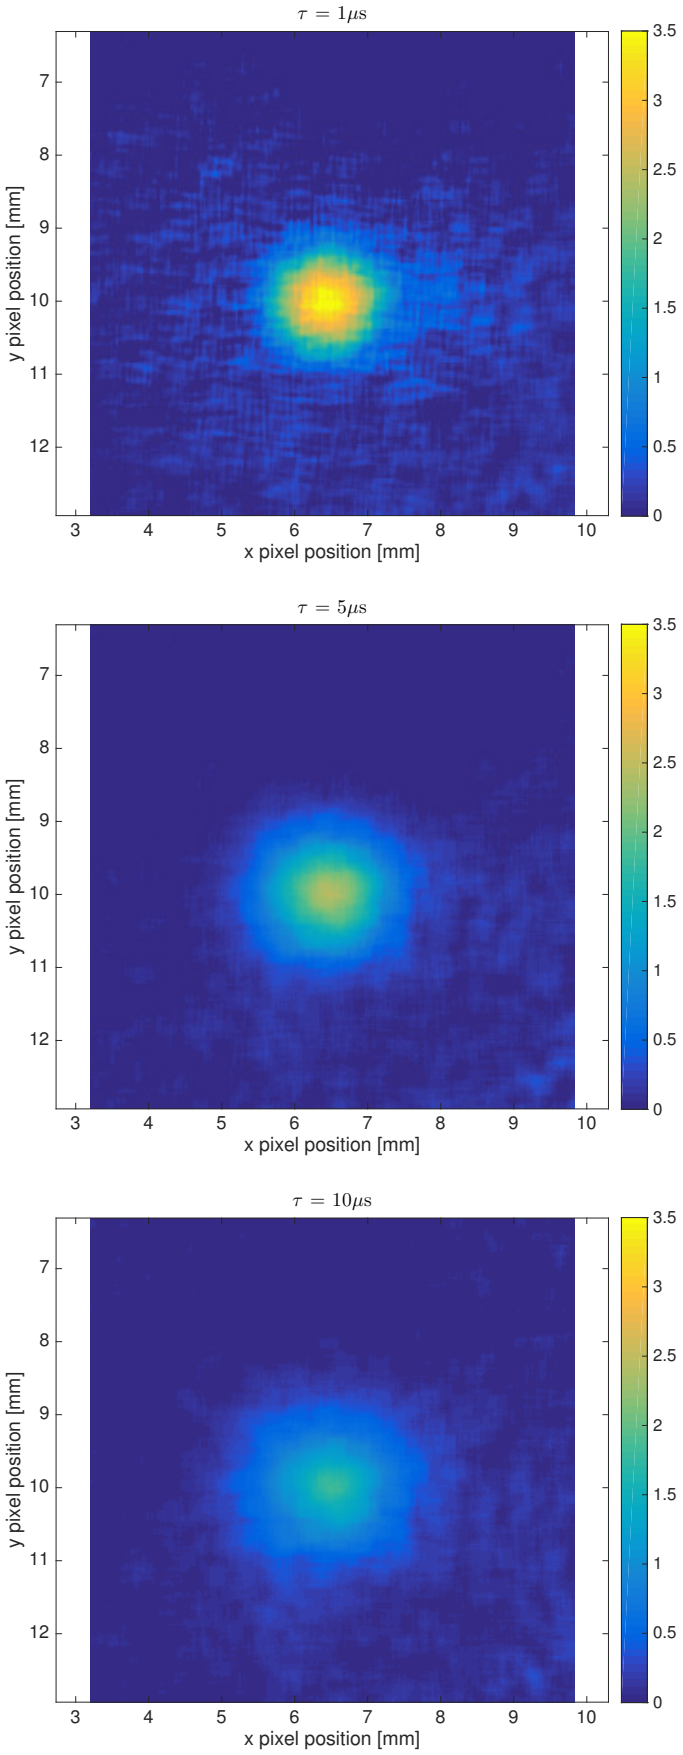


Figure 4.2: A sequence of images showing the spatial profile of Δd through the vapour cell containing neon buffer gas at 1 Torr, for delays of $1 : 10\,\mu$s

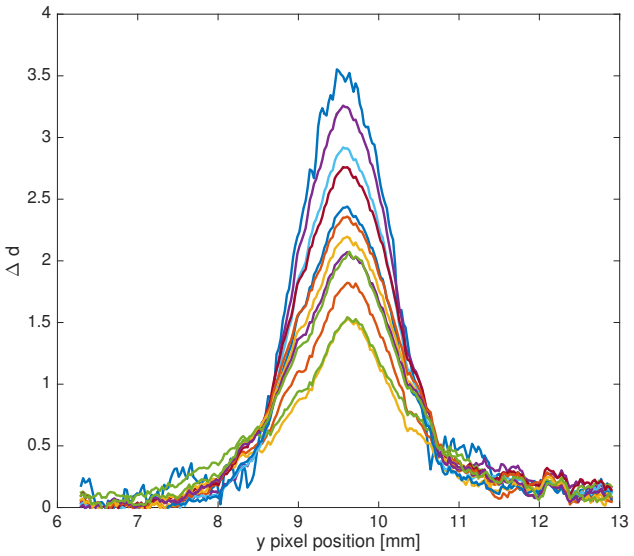


Figure 4.3: Image cross section showing variation of optical depth, for a typical delay sequence. Data from cell containing 1 Torr neon for a series of delays in the range $1:30\,\mu$s. The highest/narrowest peak corresponds to 1μs delay.

As discussed in [172], this approximation does not fully allow for experimental details which are not well described analytically, and they propose a more general method which we follow in the subsequent section.

4.2.2(b) Fourier Method

Noting the spatial asymmetries apparent in figures 4.2 and 4.3, we consider the more general approach to analyse the diffusion of the pumped atoms presented in [172] and examine the evolution different wavenumber components of the absorption profile in the Fourier domain.

We first recall the diffusion equation 4.3 for the atomic motion, from which we aim to obtain the diffusion coefficient D in the presence of different buffer gases. Experimentally we measure the transverse spatial variation in optical depth, Δd; then, transforming to the spatial Fourier domain in the transverse coordinate $\boldsymbol{\rho} = x\hat{\boldsymbol{e}}_x + y\hat{\boldsymbol{e}}_y$

$$\mathcal{F}\left\{\Delta d(\boldsymbol{\rho}, t)\right\} = \frac{1}{2\pi}\iint \mathrm{d}\boldsymbol{\rho}\, \Delta d(\boldsymbol{\rho}, t)\mathrm{e}^{-\mathrm{i}\boldsymbol{k}_\perp\cdot\boldsymbol{\rho}}, \tag{4.7}$$

for which the solution of the diffusion equation gives

$$\begin{aligned}\mathcal{F}\left\{\Delta d(\boldsymbol{\rho}, t)\right\} &= \mathcal{F}\left\{\Delta d(\boldsymbol{\rho}, 0)\right\}\, \mathrm{e}^{-\left[\gamma_0 + D|\boldsymbol{k}_\perp|^2\right]t} \\ &= \mathcal{F}\left\{\Delta d(\boldsymbol{\rho}, 0)\right\}\, \mathrm{e}^{-\gamma(k_\perp)t},\end{aligned} \tag{4.8}$$

where γ_0 is explicitly added to account for additional relaxation mechanisms. The interpretation of the above is that components of the perturbed optical depth with a given "spatial frequency" decay exponentially with rate $\gamma_0 + Dk_\perp^2$.

For a given pressure of each buffer gas, we therefore fit an exponential decay to spatial frequency components corresponding to lengthscales between the pixel size ($k_{\max} \sim (1/0.0259)\,\mathrm{mm}^{-1}$) and the probe beam size ($k_{\min} \sim (1/10)\,\mathrm{mm}^{-1}$). The decay constant, $\gamma(k_\perp)$, of the amplitude of each element in the Fourier transformed image with a given wavenumber $k_\perp = \sqrt{k_x^2 + k_y^2}$ is obtained by fitting the time dependence for a given element to a decaying exponential. The extracted values are then fitted

Species	D_0 [$\mathrm{cm^2 s^{-1}}$]
N_2	$0.223^{0.350}_{0.098}$
Ne	$0.355^{0.414}_{0.297}$
Xe	$0.055^{0.067}_{0.042}$

Table 4.2: Extracted parameters using the method presented in [172]. Superscript/subscript denote 95% confidence bounds for fit values quoted.

to a quadratic $\gamma = \gamma_0 + Dk_\perp^2$ to yield the diffusion constant D for each buffer gas at a given pressure.

The diffusion properties in the presence of a given buffer gas can be characterised by the representative parameter D_0 [173] which is related to the diffusion constant D at a given pressure, p, by:

$$D_0 = \frac{p}{760\,\mathrm{Torr}} D(p). \tag{4.9}$$

The extracted D_0 parameters using equation 4.9 are plotted in figure 4.4, and tabulated in table 4.2. The error bars are generated by performing the automated fitting routine independently over each quadrant of the image. Note that the fits have been constrained so as to pass through zero diffusion at "infinite pressure". For low diffusion coefficients (high pressures), the fits work more reliably and using the simpler Gaussian method gives good agreement; however, at higher diffusion rates, the limited bit depth of the camera and asymmetries due to the optics leads to much higher uncertainty, and the Gaussian method performs much more poorly. Whilst values for D_0 in the literature vary from source to source, the numbers obtained here are comparable to those quoted in [173].

4.2.3 Implications

In this section consider the objective of achieving a minimum coherence time[3] τ of $1\,\mu$s for a given waist size, w_0 inside a ring cavity of order M (recalling the notation introduced in section 3.3) for suppression of the anti-Stokes mode at an operating temperature where we can achieve a spin polarisation $w \geq 0.995$. Investigation of the pumping efficiency with temperature in the presence of different buffer gases was conducted in parallel by S.E.T., through which by being able to more efficiently initialise the atoms to higher temperatures significantly higher optical depths and hence memory efficiencies are achievable [142]; we will return to this in section 4.3.

The results of these investigations combined with the analysis of the diffusion properties considered here provide the bases for the choice of vapour cell used in future iterations of the cavity memory experiment.

In these parallel investigations, we considered two alternatives, nitrogen and xenon, to neon as buffer gases. With the additional benefit of quenching possible with nitrogen, much higher optical depths can be achieved, we therefore choose[4] to move towards substituting Ne with N_2.

With the above objective for the coherence time in mind, we estimate the expected *diffusion-limited* memory lifetime using a similar analysis as in the previous section. By considering a Gaussian optics

[3] Bearing in mind that, as discussed in section 4.1.2, with our magnetic shielding the timescale for magnetic dephasing was shown to have a lower limit $\gtrsim \mu$s.

[4] Note also that one might consider using a mixture of both N_2 and Xe buffer gases, attempting to simultaneously exploit the slower diffusion rate with xenon and the quenching properties of nitrogen. However, this mixture would impede the ability to quench. This is because the pressure linewidth of Xe is approximately double that of N_2, and thus the probability of non-radiative decay through a collision with nitrogen would be reduced due to the decreased ratio of the contribution from non-radiative decay channels to the excited state lifetime.

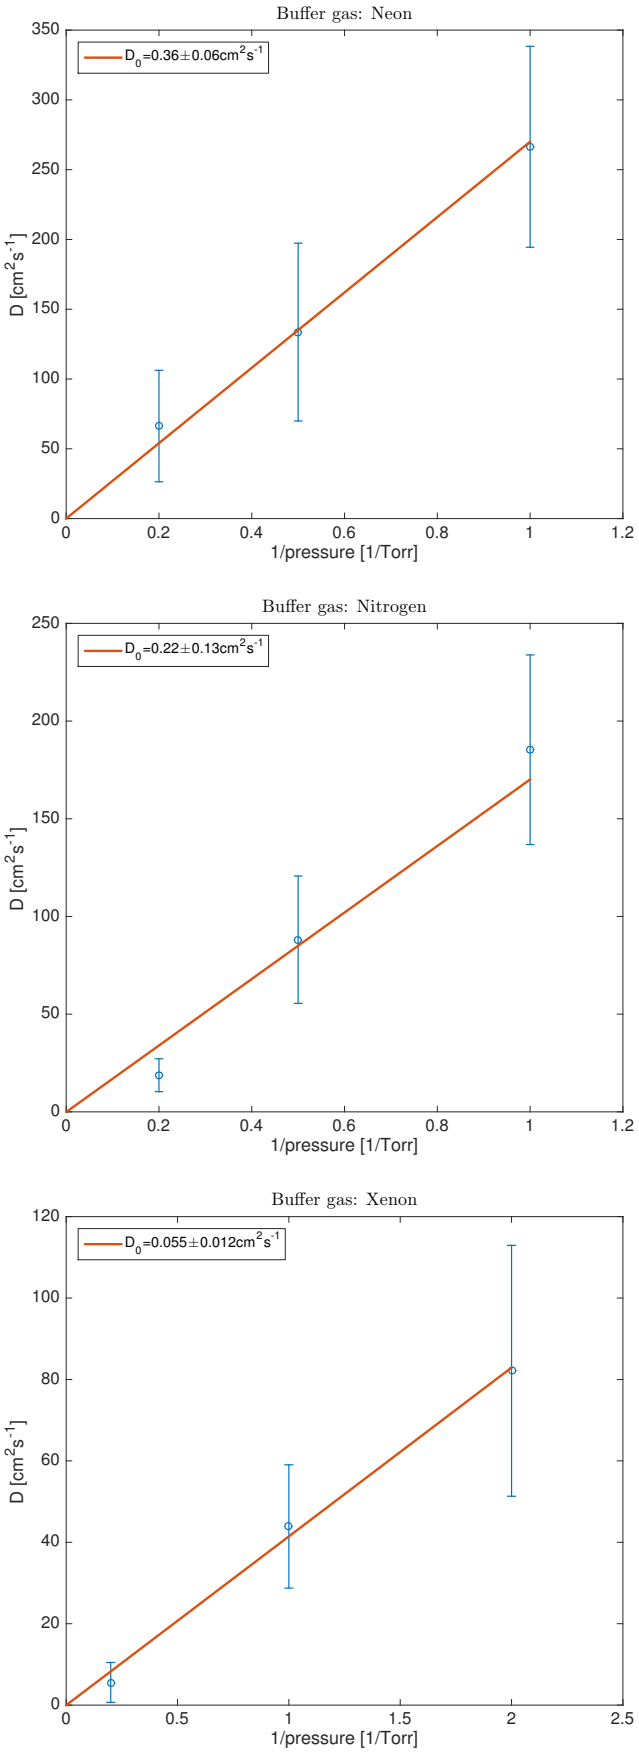


Figure 4.4: Extraction of D_0 for three buffer gas species.

M	L [mm]	R [mm]	w_0 [μm]
2	40.8	40	73.7
1	24.5	25	58.2
0	8.15	10	36.5

(a) Beam waist for small R mirrors

	w_0 [μm]	R [m]	
		1	3
M	2	195.8	258.6
	1	172.7	227.7
	0	131.5	173.2

(b) Beam waist for large R mirrors.

Table 4.3: (a) Target beam waist for different order cavities given our chosen cavity length and the optics available, for maximising intracavity Rabi frequency (small R) ; (b) Required waist sizes for large radius of curvature mirrors, which could be used for any of the cavity orders.

model for the ring cavity (neglecting the atomic contribution to the cavity response) with $\mathrm{FSR}_M^{(0)}$, we can obtain an estimate of the waist size required given the mirrors that can be manufactured for a given target cavity length.

We consider the propagation of a Gaussian mode around the cavity with the "$ABCD$" matrix [174] from the centre of the cavity

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} 1 & L/2 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ -R/2 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & L/2 \\ 0 & 1 \end{pmatrix} \tag{4.10}$$

for a single roundtrip cavity length L with a concave mirror, radius of curvature R. Requiring that the q parameter of the beam is unchanged after one roundtrip and the waist is located at the centre of symmetry of the cavity we must have that $R > L/2$. With standard components meeting this condition we arrive at the beam waists tabulated in 4.3a.

If these do not result in sufficient enhancement in lifetime, we may move to long radius of curvature mirrors at the expense of attainable intracavity Rabi frequency in the control field. The larger radius mirrors in this case could be applicable to any of the cavity orders.

For an interaction region of radius r (related to the beam waist w_0 by some factor) and length L, outside of which regions the atoms are "lost", the loss rate due to diffusion can be expressed [168]:

$$\Gamma_{\text{dif.}} = \left[\left(\frac{\pi}{L}\right)^2 + \left(\frac{\mu}{r}\right)^2\right] \times D(p), \tag{4.11}$$

where $\mu = 2.405$ is the first zero of the zero$^{\text{th}}$ order Bessel function, and $D(p)$ is obtained from D_0 via equation 4.9.

To obtain an estimate of the coherence time we can consider a simple picture using the Gaussian approximation above, whereby we consider the overlap of the beam area with the atoms at some later time using the expression in equation 4.6 such that the waist of the distribution evolves as

$$\sigma(t) = \sigma_0\sqrt{1+\tau}, \tag{4.12}$$

where $\tau = 4Dt/\sigma_0^2$ is a dimensionless time parameter, expressed in terms of the initial beam width σ_0 and the diffusion constant D.

If at dimensionless time $\tau = 0$ the atoms are initialised with a Gaussian spatial distribution, then the number of atoms which have been successfully initialised, per unit length in an anulus at some

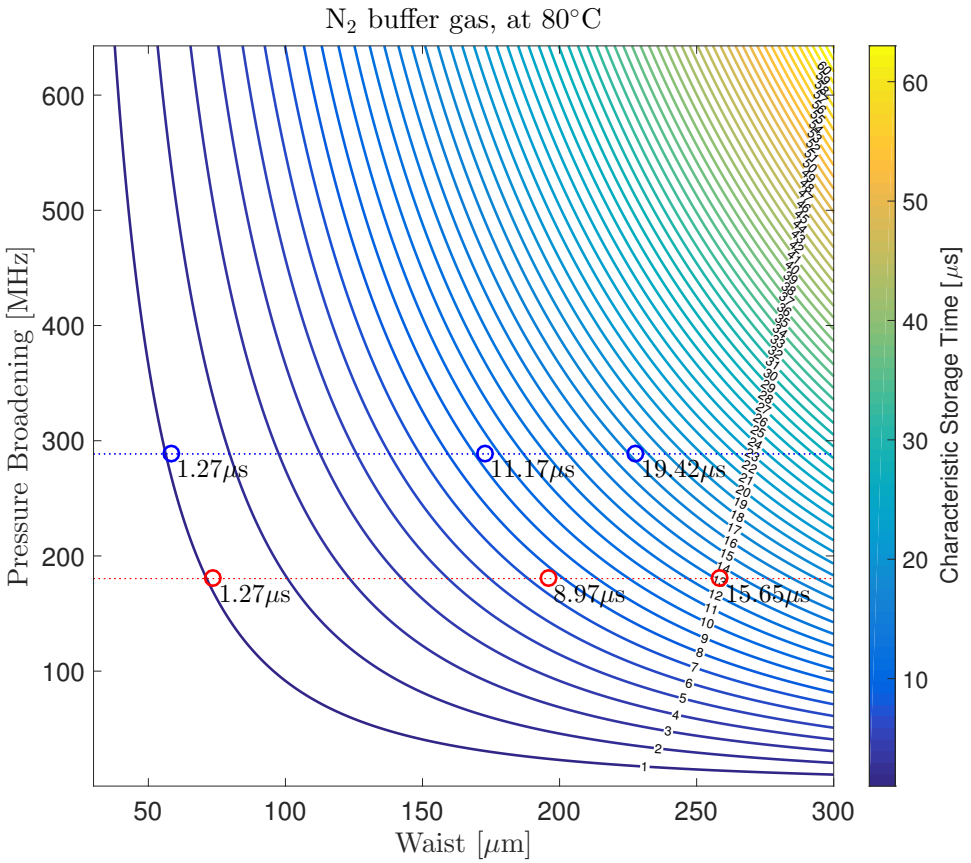


Figure 4.5: Predicted dependence of diffusion limited lifetime assuming nitrogen buffer gas operating at 80°C. The red and blue lines indicate where we aim to operate with the new components for the experiment (red: $M = 2$ and blue: $M = 1$).

radius ρ is given by

$$\mathrm{d}n(\rho,\tau) = n(\rho,\tau) \times 2\pi\rho\,\mathrm{d}\rho \tag{4.13}$$

$$= \left\{ \frac{1}{2\pi\sigma(\tau)^2} \exp\left[-\frac{1}{2}\left(\frac{\rho}{\sigma(\tau)} \right)^2 \right] \right\} \times 2\pi\rho\,\mathrm{d}\rho, \tag{4.14}$$

such that $\int \mathrm{d}n(\rho) = 1$. When the probe is then applied some time, τ, later, the atoms will have begun to diffuse out of the interaction region, and we can estimate the loss due to diffusion simply by the overlap between the initial Gaussian region and the time evolved distribution.

The resulting e^{-2} timescale is then $\tau = 2(\mathrm{e}^2 - 1)$. This estimate quantitatively corresponds very well with that arrived at using equation 4.11 if r is taken as ~ 1.5 times the beam waist. The trade off between beam waist and pressure broadening with respect to the predicted lifetime due to diffusion is plotted in figure 4.5.

Based upon the results in [142], which I summarise in the following section 4.3, we assume that we can now efficiently pump (polarisation 99.5%) the ensemble up to temperatures ~ 80°C by pumping on the D_1 line with ~ 10 Torr nitrogen buffer gas. In this case, the projected lifetime can be estimated for an $M = 2$ cell (10 mm) with 10 Torr nitrogen and an $M = 1$ cell (6 mm) with 16 Torr nitrogen, which we have orderered for the next planned iterations of the experiment are shown in table 4.4. Note that the two small R mirrors are estimated to achieve our target lifetime $> 1\,\mu$s if the same level of magnetic shielding is used as in the bulk experiments; however, if necessary we may move to the larger $R \gtrsim 1$ m mirrors, sacrificing the peak intracavity Rabi frequency.

order	M	2	1
cell length	[mm]	10	6
Pressure	[Torr]	10	16
(minimum R) beam waist	[μm]	73.7	58.2
Projected lifetime	[μs]	**1.27**	**1.27**
($R = 1$m) beam waist	[μm]	195.8	172.7
Projected lifetime	[μs]	**8.97**	**11.2**
($R = 3$m) beam waist	[μm]	258.6	227.7
Projected lifetime	[μs]	**15.6**	**19.6**

Table 4.4: Estimated experimental parameters, and projected coherence times achievable for nitrogen buffer gas operating at 81°C, assuming diffusion-limited decoherence.

If we are able to achieve sufficiently long lifetimes such that $\tau_{\text{dif.}} \sim \tau_{\text{mag}}$ with this step, further efforts will be needed to improve the magnetic shielding in order to increase the lifetime further.

4.3 Buffer Gas (2): Quenching

The Raman memory protocol, as many others, requires that we efficiently initialise the ensemble,[5] and, as the memory efficiency in the Raman memory scheme is determined by the Raman coupling constant, we seek to maximise the caesium number density (via increased temperature). However, elevated number densities leads to radiation trapping at high density, refer again to figure 3.3b whereupon after resonant excitation the atom will decay via spontaneous emission to one of the ground states. At low atomic densities, continual depletion of one of the ground states via the pumping laser will eventually result in a spin polarised ensemble; at high atomic densities however, this emitted photon has an increased probability of exciting another atom in the ensemble. This process can be parametrised in terms of a multiplicity factor $\mathcal{M}$ which gives the average number of times a photon is reabsorbed before it exits the atomic ensemble [168].

A known means to impede this is via "quenching" [173], where by introducing a molecular buffer gas, which exhibits a mechanical degree of freedom near in energy to the atomic transition, therefore enables the non-radiative decay of the atoms from the excited state. Optical pumping efficiencies as a function of temperature were therefore measured for a number of buffer gas configurations.

Historically, we have used neon buffer gas with optical pumping on the D_2 line, however the result of the investigation led by S.E.T. [142] has shown that with 10 Torr of N_2 and by pumping on the D_1 line we can now achieve polarisations of over 99.5% up to 90°C. This increase in temperature from 70°C increases the optical depth by a factor of ~ 4 [142]. These improvements are possible since the quenching cross section (the probability of a quenching collision occuring) is greater for the D_1 transition, and additionally excited state lifetime of $6^2P_{1/2}$ is longer than for the $6^2P_{3/2}$ so there is a higher probability of nonradiative decay occuring during the excited state lifetime.

The ability to effectively optically pump at these increased optical depth has enabled us to operate the memory in a much higher coupling regime, to the extent that FWM gain becomes the dominant interaction, such that the efficiency defined as $\sim N_{\text{out}}/N_{\text{in}}$, may be greater than unity due to gain. In fact, a signal amplification of 20 dB in this system has been observed [142], although this is not

[5] This is not a universal problem – a new memory protocol "ORCA" developed more recently by K.T.K in our group requires no initialisation [115].

useful for a quantum memory,[6] this dramatic increase does demonstrate that we are now able to operate in a regime with significantly enhanced coupling strength. With these elevated pumping efficiencies and increased optical density, we estimate the efficiency of the *noise suppressed* cavity Raman memory implemented in the previous chapter would be approximately doubled [142] without any further optimisation.

4.4 Numerical Parameter Search

Following on from the proposal, theory and first experimental demonstration of the cavity Raman memory, presented in the previous chapter, the next phase of the project is to work towards interfacing this with single photon source and look at photon statistics. To lay the ground work for this, and to supplement the analysis of projected lifetimes which may be achieved given in the previous section, we now briefly discuss some other characteristics of the memory performance using the theoretical framework we presented in the previous chapter (and references [147, 154]).

To illustrate how these considerations may be used for the design of subsequent generations of the memory experiment we highlight some key results of the theory for the cavity memory from the previous chapter, and then discuss how these may be incorporated with the numerical models of the "atom plus cavity" system.

4.4.1 Cavity Raman Memory

With the introduction of the cavity it is a known result that the Raman memory efficiency depends upon the cooperativity $\mathcal{C}$ at the signal frequency [138], where (as in equation 3.35) the cooperativity incorporates the cavity enhanced optical depth of the ensemble, and can be written as

$$\mathcal{C} = \frac{d}{1 - \mu_{\mathrm{s}} \mathrm{e}^{\mathrm{i}\phi_{\mathrm{s}}}}, \tag{4.15}$$

where d is the roundtrip optical depth of the atomic ensemble, $\mu_{\mathrm{s,a}} \in \Re$ is the full roundtrip amplitude transmission of the cavity including reflectivities, atomic absorption and any other losses, and $\phi_{\mathrm{s,a}}$ the roundtrip phase accumulated at the signal "s" or anti-Stokes "a" frequencies. This is closely related to the cavity transmission parameter, which we define

$$\chi_{\mathrm{s,a}} = \frac{1 - r^2}{1 - \mu_{\mathrm{s,a}} \mathrm{e}^{\mathrm{i}\phi_{\mathrm{s,a}}}}, \tag{4.16}$$

which will also become relevant later, where r is the reflectivity of the input coupler.

We then define the memory coupling parameters for the signal/anti-Stokes fields

$$C_{\mathrm{s,a}} = \frac{\sqrt{\mathcal{C}\gamma W}}{\Gamma_{\mathrm{s,a}}}, \tag{4.17}$$

where γ is the homogeneous linewidth of the excited state, $W = \int_{-\infty}^{\infty} \mathrm{d}t' |\Omega(t')|^2$ is proportional to the control pulse energy, and $\Gamma_{\mathrm{s,a}} = \gamma - \mathrm{i}\Delta_{\mathrm{s,a}}$ is the complex detuning. Note that in this definition both couplings are expressed in terms of the cooperativity *at the signal frequency*.

We then recall that in the dynamical equations, the Raman coupling for the anti-Stokes, C_{a} always

[6] This could have applications in the generation of strongly squeezed light via FWM in vapour e.g. [175].

appeared with a prefactor x with respect to the Stokes (signal) coupling C_s where

$$x = \frac{1 - \mu_\mathrm{s} \mathrm{e}^{+\mathrm{i}\phi_\mathrm{s}}}{1 - \mu_\mathrm{a} \mathrm{e}^{-\mathrm{i}\phi_\mathrm{a}}}. \tag{4.18}$$

Appearing in this fashion, x therefore encapsulates the noise suppression due to the cavity, which we aim to minimise in order to reduce the four-wave mixing noise.

As one would intuitively expect, we therefore seek a "double resonance" condition whereby $\phi_\mathrm{s} = 0 \bmod 2\pi$ and $\phi_\mathrm{a} = \pi \bmod 2\pi$, i.e. that the signal is resonant with the cavity, and the anti-Stokes anti-resonant. If this is met, it is apparent that the achievable noise suppression scales with the finesse of the cavity. For a given cavity order and target free spectral range, there is therefore a trade off between the noise suppression attainable and the acceptance bandwidth (or equivalently the speed of operation). Note that this model assumes adiabatic operation: that we are in the "bad cavity" limit where external fields are spectrally much narrower than the decay rates of the cavity, thus the bandwidths used should be smaller than the acceptance bandwidth by some margin.

Phenomenologically, we can express the number of photons recalled from the memory in the retrieval time-bin in the form

$$N_\mathrm{out} = \eta N_\mathrm{in} + N_\mathrm{noise}, \tag{4.19}$$

where η is then the memory efficiency, and $\eta N_\mathrm{in}/N_\mathrm{noise}$ is the signal-to-noise ratio (SNR), and N_in is the average number of photons per write pulse. For a heralded photon source, N_in can be interpreted as the heralding efficiency.[7] We note that framed in this way, the resulting efficiencies may be greater than unity, since four-wave mixing gain can boost the efficiency whilst introducing noise in the amplification process. Another key figure of merit for the operation of a quantum memory is the ability to preserve the statistics of the stored states, which we here quantify by computing the $g^{(2)}$ autocorrelation function for the fields emerging from the memory. The expressions for η, SNR and $g^{(2)}$, and some limiting cases, were derived in section 3.2.

We must also bear in mind the control field Rabi frequency, which, without perfect mode matching between signal and control, causes the efficiency to roll off at higher pulse energies due to the a.c. Stark shift moving the system out of two-photon resonance. This can be circumvented by shaping the control field [141], in which case, the peak memory efficiency attainable is then determined by the cooperativity of the cavity [138].

Note that the deriviation of the performance metrics was based on a number of approximations[8] in order to ultimately obtain insightful equations of motion. Given the sensitivity of these results to the dispersion and intracavity loss, in the following section we interface these with a more complete model of the cavity response evaluated using the full hyperfine structure of the atoms and physically reasonable component properties.[9] We investigate the trade off that may be afforded between varying temperature, detuning and control field power.

Having noted the asymptotic limit of the efficiency determined by the intracavity loss, the cavity configurations explored here are all *single port*, i.e. we operate the cavity in reflection as in figure 3.2; which additionally prevents degradation of the efficiency by read-out from multiple ports.

[7] More explicitly, the probability of having a photon at the input of the memory.
[8] Such as coarse graining to a three level system, and Taylor expansion to capture roundtrip evolution.
[9] For instance imperfect anti-reflection (AR) coating on the cell windows.

4.4.2 Numerical Implementation

To numerically explore this wide parameter space, we mimic an experimental procedure: whereby for some cavity configuration – mirror reflectivities, vapour cell temperature and estimate of the cell length, and then "lock" the cavity length[10] to resonance with the signal at some chosen detuning.

We then simulate the full cavity response using the infrastructure in section 3.3. From the cavity response, we can then extract the necessary parameters $\chi_{s,a}$, $C_{s,a}$ and x to then obtain the memory performance metrics derived in section 3.2.

An example of this process is shown in figure 4.6 for a single port ring cavity (for (a) $M = 2$, and (b) $M = 1$); where we consider starting from 70°C at 15.2 GHz detuning (for historic reasons comparing with [118]) and consider varying the temperature of the vapour cell, and adjusting the detuning *so as to keep the loss at the signal frequency constant.* i.e. as the temperature (and hence optical depth) is increased, we move to a greater detuning to compensate the increased linear absorption of the atoms. Note that in this process, as with the experiment, the anti-Stokes is not perfectly anti-resonant with the cavity (since the dispersion from the atoms serves to reduce the FSR with respect to the empty cavity used to initially estimate the length).

A possible optimisation procedure here is to identify a target value for x given an acceptable trade off between the bandwidth of the signals that may be stored and achievable noise suppression. From this we can obtain an estimate of the roundtrip loss, which, in combination with the performance of the cavity components dictates a level of linear absorption permissible within the atomic ensemble. The linear absorption $\alpha \sim d(\gamma/\Delta)^2$, so depends on the temperature (through d) and detuning.[11] These degrees of freedom can then be optimised upon whilst considering the predicted memory efficiency, signal to noise ratio and $g^{(2)}$ of the retrieved state.

In this case, since the dispersion drops off at a slower rate than the absorption with detuning, maintaining a constant loss whilst increasing the temperature/optical depth (and hence necessarily increasing detuning) will result in an *increase in atomic dispersion.* In attempting to apply the theoretical analysis within the [147] (`arXiv v.1`) with the simulated cavity responses, it was noted that some approximations used may lose validity at the elevated pressure broadening which were used in the experiment of the previous chapter, and those we considered in the previous section in order to facilitate achieving the desired lifetimes.

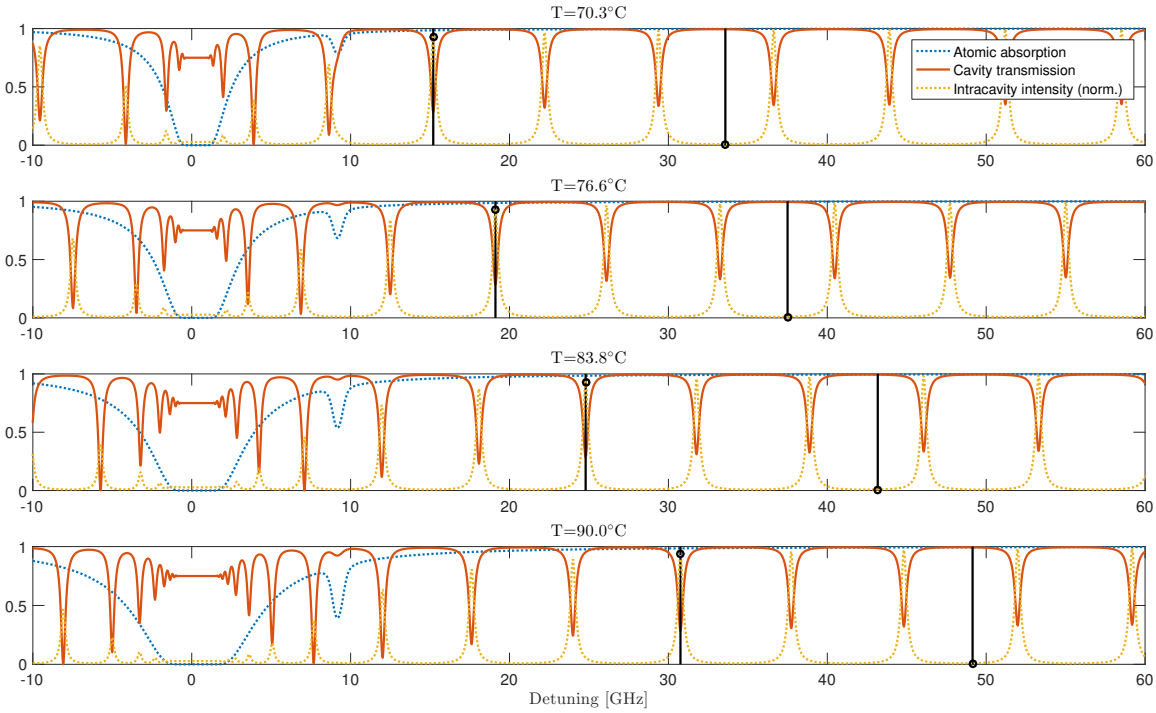
This means that the assumption that the atomic contribution to absorption and dispersion are *small* and may be approximated by a Taylor expansion becomes less valid. Together J.N. and I modified the analysis, to handle the roundtrip phase more carefully[12] – as is outlined in the derivation of the previous chapter and as now appears in [147]. Nevertheless, in the future a more full analysis of the signal, spin wave and anti-Stokes evolution may be necessary, for example via computation of the Green's function, c.f. [176].

For demonstration purposes, these results were obtained using the original model and we consider an example of a low loss $M = 0$ cavity, with negligible pressure broadening where this model may be used more reliably. We also assume perfect mode matching (as defined by equation 3.77), which can be achieved by shaping the control pulse. For comparison, we otherwise take the signal and control temporal profiles as identical, and the effiency rolls off at higher control powers and the efficiency may

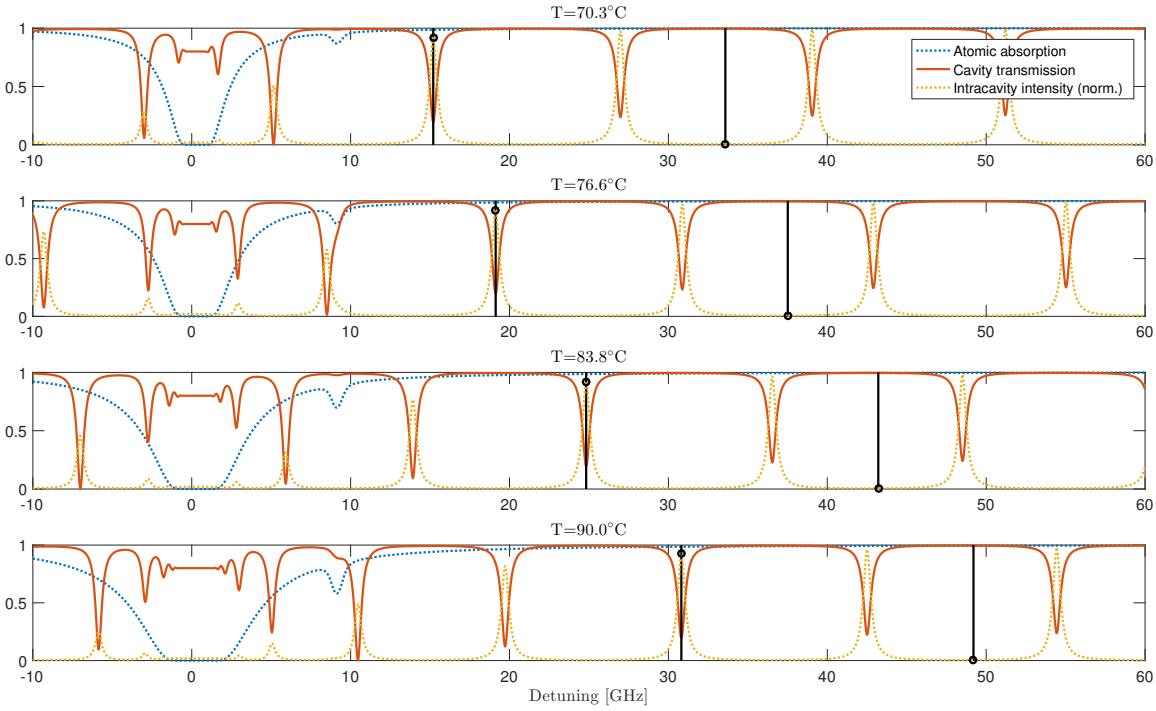
[10] i.e. by modifying the path length outside of the cell.

[11] Note also that, regardless of the cavity, the noise suppression also depends on the ratio of detunings Γ_s/Γ_a, i.e. the intrinsic suppression is reduced at larger detunings.

[12] Resulting in the modification of some of the final variable definitions, but qualitatively similar conclusions.



(a) Second order cavity ($r_{\mathrm{in}} = 0.75$, 10 Torr N_2).



(b) First order cavity, ($r_{\mathrm{in}} = 0.80$, 16 Torr N_2)

Figure 4.6: Cavity response with laser detuning from the populated atomic transition, showing the absorption spectrum of the atoms (blue, dotted), the cavity transmission at the output port (orange), and the intracavity field intensity (yellow, dotted). Simulated for second (a) and first (b) order cavity, where the detuning is adjusted as the temperature is varied so as to keep the intracavity loss at the signal frequency constant with respect to operation with 15.2 GHz detuning at 70°C as used in the bulk experiments. The black lines and circles indicate the positions of the signal and anti-Stokes frequencies. Note that the response, using the trial length estimated for an empty cavity, applying a small correction to adjust the lock position into resonance with the signal is therefore not perfectly anti-resonant with the anti-Stokes since the additional dispersion from the atoms reduces the FSR.

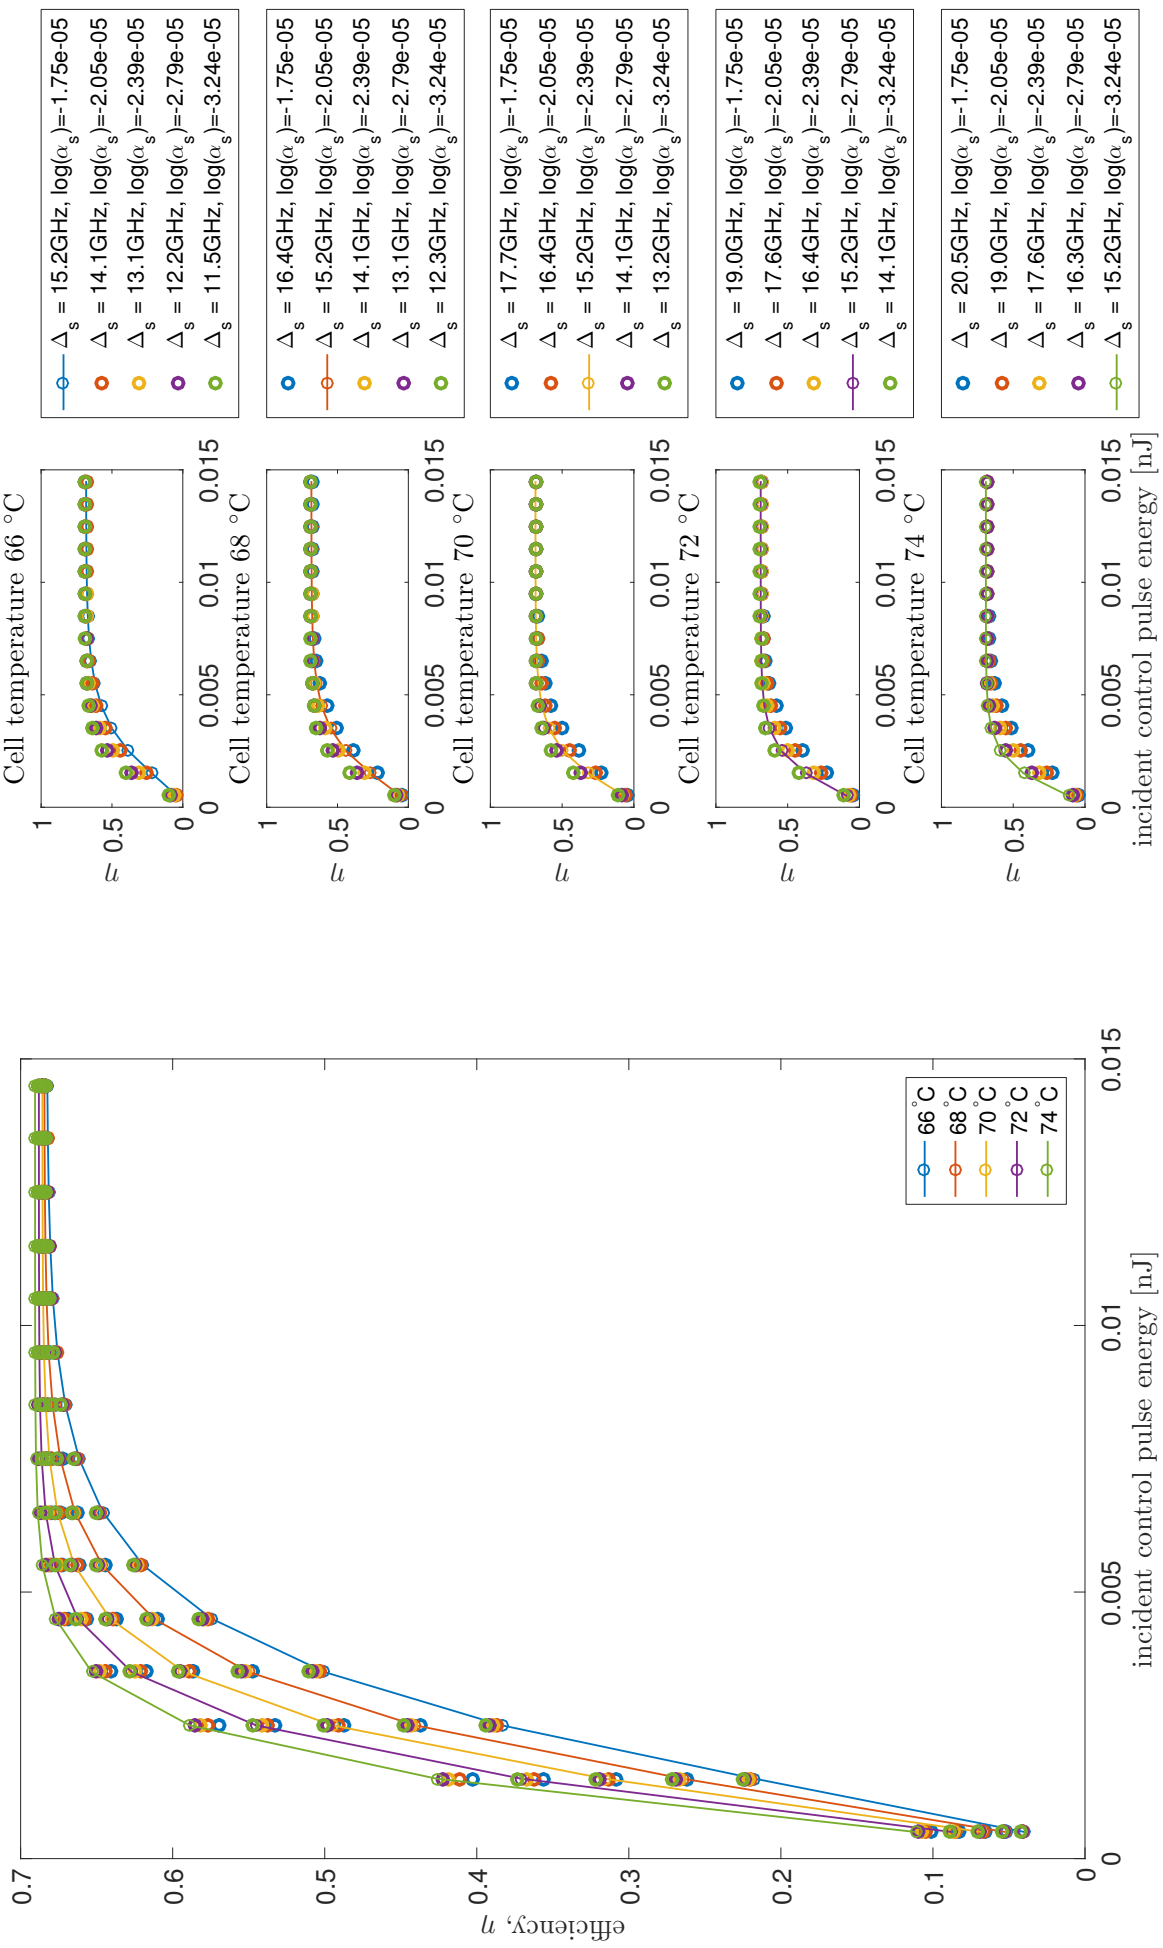
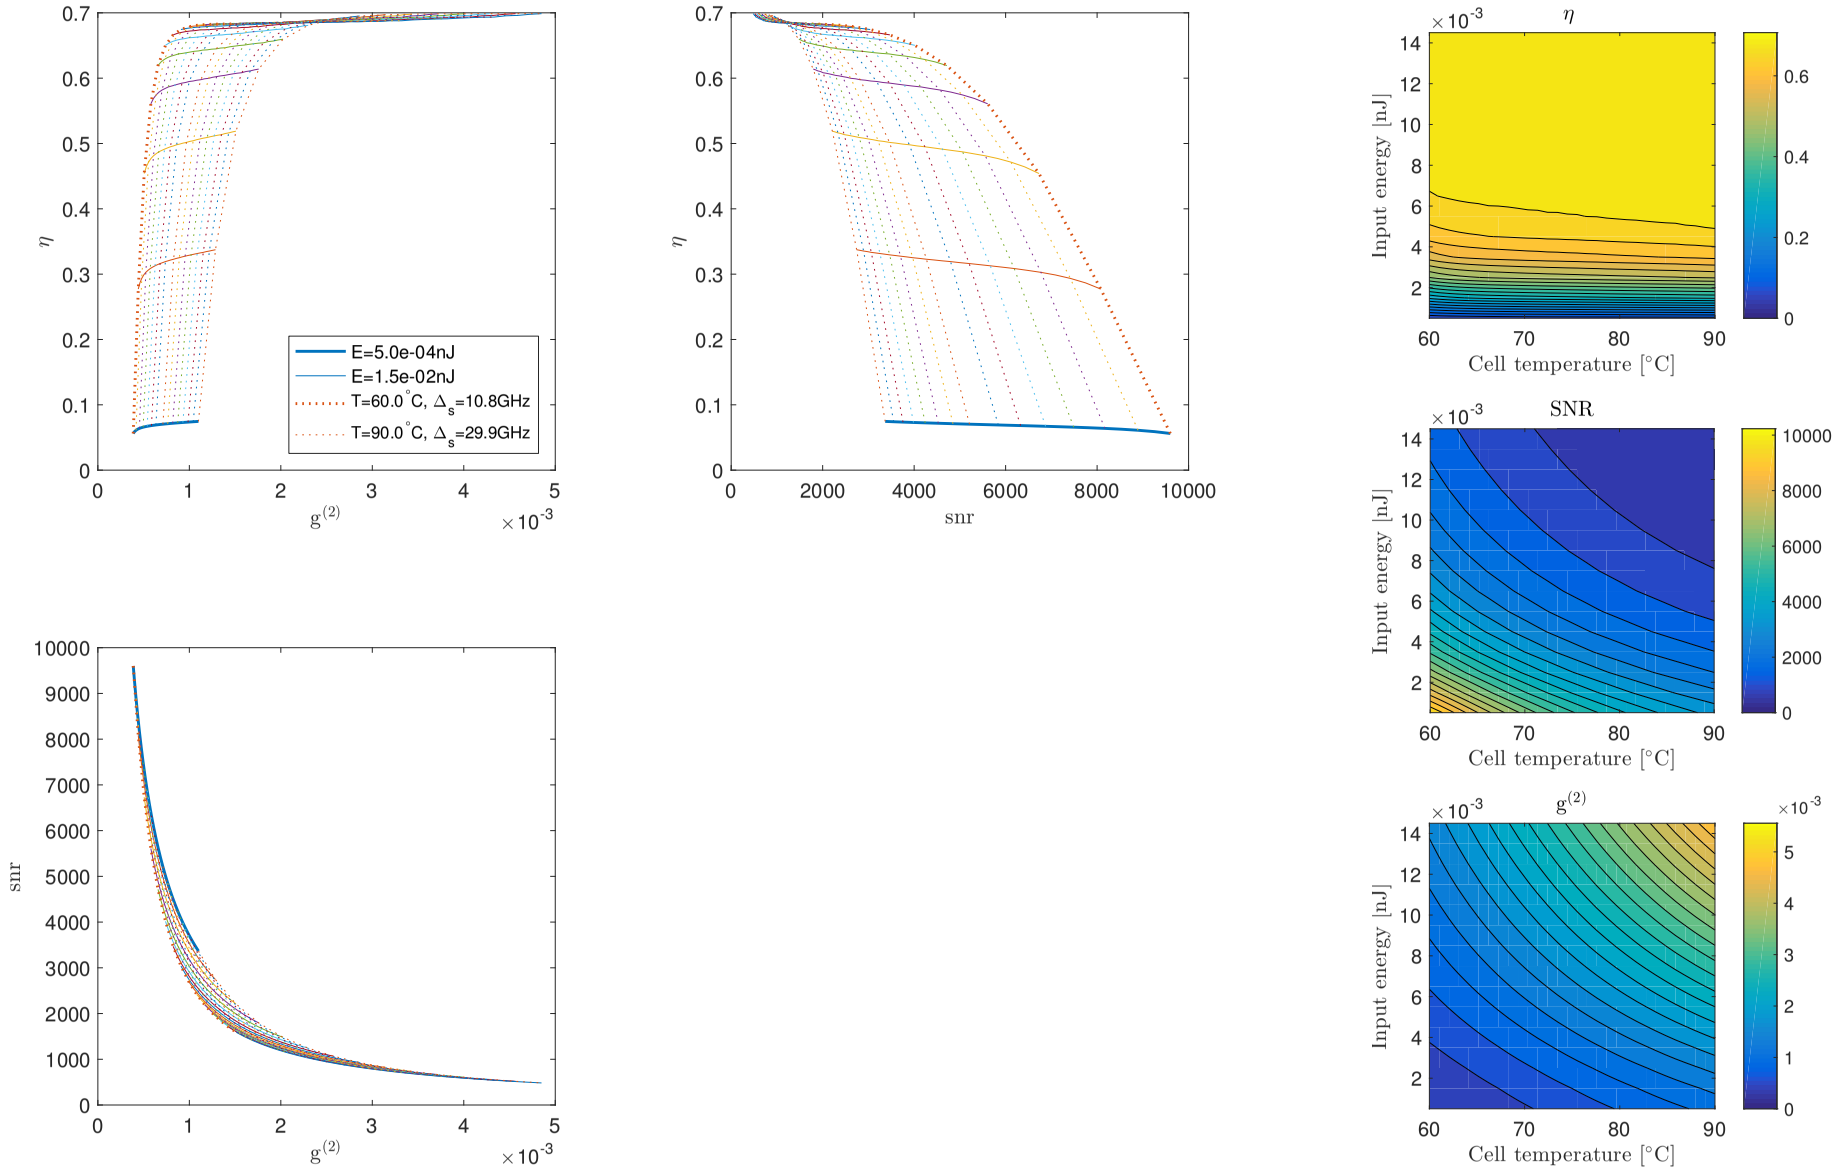
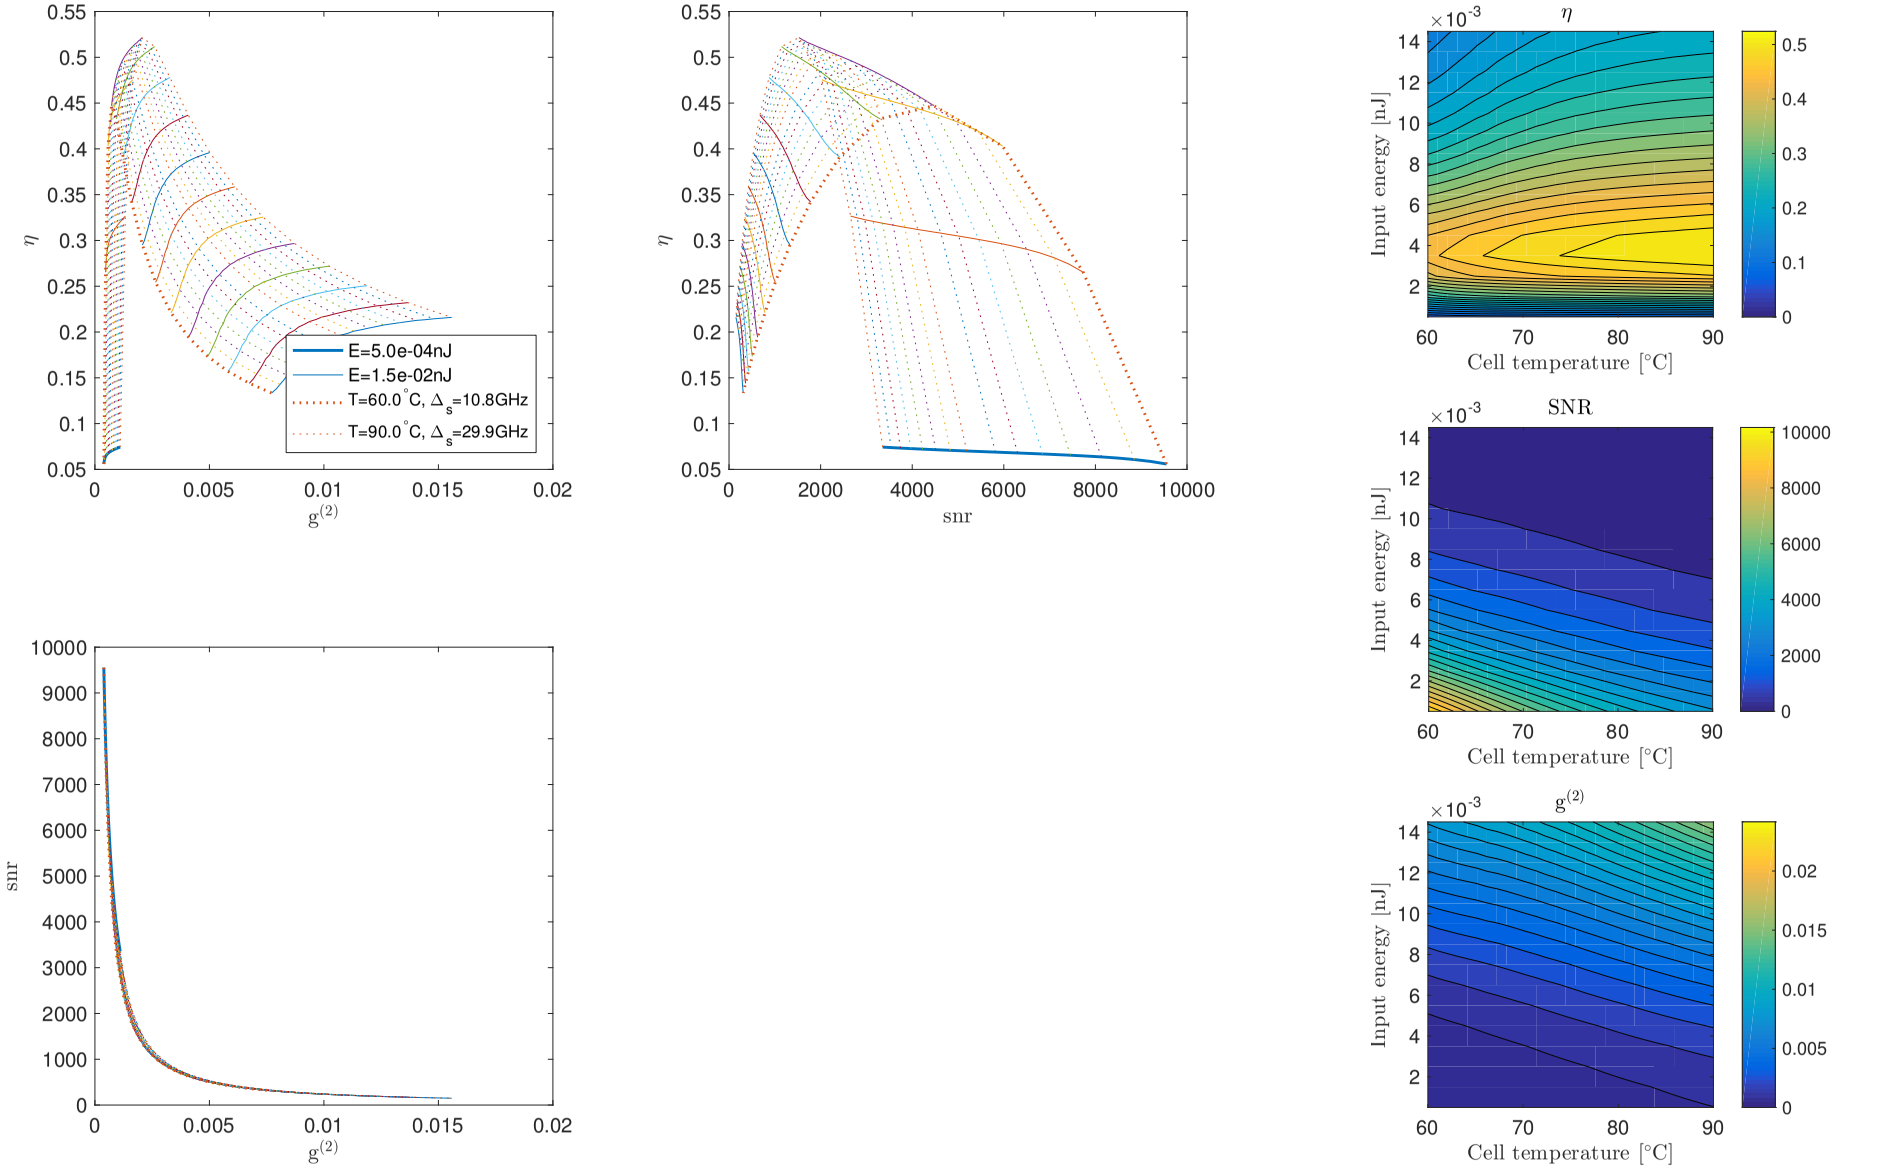


Figure 4.7: Variation of predicted memory efficiency for a low loss $M = 0$ cavity, with low pressure broadening. Series of the same colour have the same loss at the signal frequency. Left: each "bundle" of points are at the same temperature, but with different detuning (and hence loss), the solid lines (giving the temperature of each set) are all generated at the same detuning (15.2 GHz). Right: each temperature series has been seperated out. Increasing the temperature causes the efficiency to rise faster with control field Rabi frequency, whilst varying the loss determines the maximum efficiency at which the memory saturates. N.B. I use nJ here for pulse energies comparison with previous implementations requiring ~ 1 nJ in the bulk [118] and the first cavity demonstration [153].



(a) $\mu = 1$, perfect mode matching

Figure 4.8: Projection of three simulated parameters to optimise, for the $M = 0$ cavity considered in the text, with a heralding efficiency of 50%. We vary the cell temperature and control field power, *adjust the detuning to maintain constant atomic absorption*. Dotted lines indicate constant temperature surfaces, and solid lines constant control pulse energy. The thick lines indicate the minimum of each of these within the range considered. (a) shows the case of perfect mode matching, $\mu = 1$.



(b) Imperfect mode matching

Figure 4.8: (cont.) (b) shows the corresponding plots with imperfect mode matching – assuming identical control field and signal temporal mode. In both cases, increasing the control field power initially increases the efficiency, but continually increases the noise contribution. With perfect mode matching, increasing the temperature reduces the control field power needed to achieve a given efficiency, and that necessary to reach saturation. Without compensation for the a.c. Stark shift, increasing the temperature enables the maximum efficiency to be increased before the Stark shift roll-off.

$10^3 \times \lvert x\rvert^2$		r		
		0.75	0.80	0.95
	2	9.2	6.3	1.3
M	1	8.8	6.1	1.2

(a) Suppression factor

$10^2 \times \frac{\lvert\chi\rvert^2}{4}$		r		
		0.75	0.80	0.95
	2	53.5	48.4	14.0
M	1	53.7	48.6	14.1

(b) Maximum efficiency

Table 4.5: (a) Resulting suppression factor, x, for different reflectivity mirrors used in different cavity orders; (b) χ parameter determines the maximum achievable efficiency, as it encodes the intracavity loss. See also figure 4.9.

not saturate. The exploratory results in this case are presented in figures 4.7 and 4.8 with $r_{\text{in}} = 0.95$ and a roundtrip transmission from optics, 99.5%. The control field energy evaluates the intracavity Rabi frequency assuming the tightly focussed beam waist given in table 4.3a, where we additionally calculate the cavity enhancement of the control field where the cavity is simultaneously resonant with the control and signal.

Whilst the original simulations could not confidently extend to the operation regime for the experimental results of the previous chapter – due to the high optical depths and large pressure broadening, these results do suggest that the ability to move to higher temperatures and increased detuning may be able to offer enhancements, particularly before optimal mode matching via pulse shaping the control field can be implemented. Nevertheless, from the results presented here, one concludes that the noise suppression factor x and transmission parameter χ for a given "cavity plus atoms" configuration are useful parameters determining the memory performance, even if the functional form of the dependence for the resulting efficiencies and output statistics did not capture the full physics.

The x and χ values for our cavity components are given in tables 4.5a and 4.5b, respectively, evaluated at 15.2 GHz detuning and a temperature 70°C, for the cells in table 4.4. We assume an intrinsic loss from the remaining cavity components giving a roundtrip amplitude transmission of $\sqrt{0.93}$ determined from measurement, and assume negligible pressure broadening. Note the intuitive trend of increasing cavity order, or reflectivity, results in smaller x, and hence greater suppression of the four-wave mixing, which can be afforded at the expense, say, of the acceptance bandwidth of the cavity. The asymptotic efficiency is given by $\eta_{\text{max}} = |\chi|^2/4$ or $\sim ((1-r^2)/(1-\mu))^2/4$ on resonance with the signal. Once the intrinsic intracavity loss has been minimised (i.e. maximising μ), yet higher efficiencies may be achieved by *increasing* the in-coupler reflectivity, although this is at the expense of acceptance bandwidth.

To clarify this, and complement the values in tables 4.5, figure 4.9 displays the asymptotic memory efficiency as the incoupler reflectivity and roundtrip transmission[13] are scanned. Here, a "roundtrip transmission" of unity indicates *no intracavity loss*. The maximally achievable efficiency at a given reflectivity – corresponding to critical coupling – varies as the intracavity loss is scanned. If the intracavity loss can be reduced sufficiently then the optimal efficiency can be increased towards unity. The figures on the right then show the inverse relationship: showing the same efficiency for each of the three mirrors considered, as the intracavity loss is scanned. The lower right figure shows the cavity linewidth at the signal frequency, via equation 3.15, estimated for the $M = 0$ cavity length.[14]

Finally, the modified theoretical description (as appears in the previous chapter) is more suitable to describe a wider range of operating conditions. Figure 4.10 plots the projected efficiency and

[13] Including all *other* intracavity losses aside from the input mirror: linear atomic absorption or scattering from optics.

[14] For a given loss, the acceptance bandwidth is decreased as the roundtrip time is increased – i.e. for longer cavities it will take more time for the intracavity field to decay.

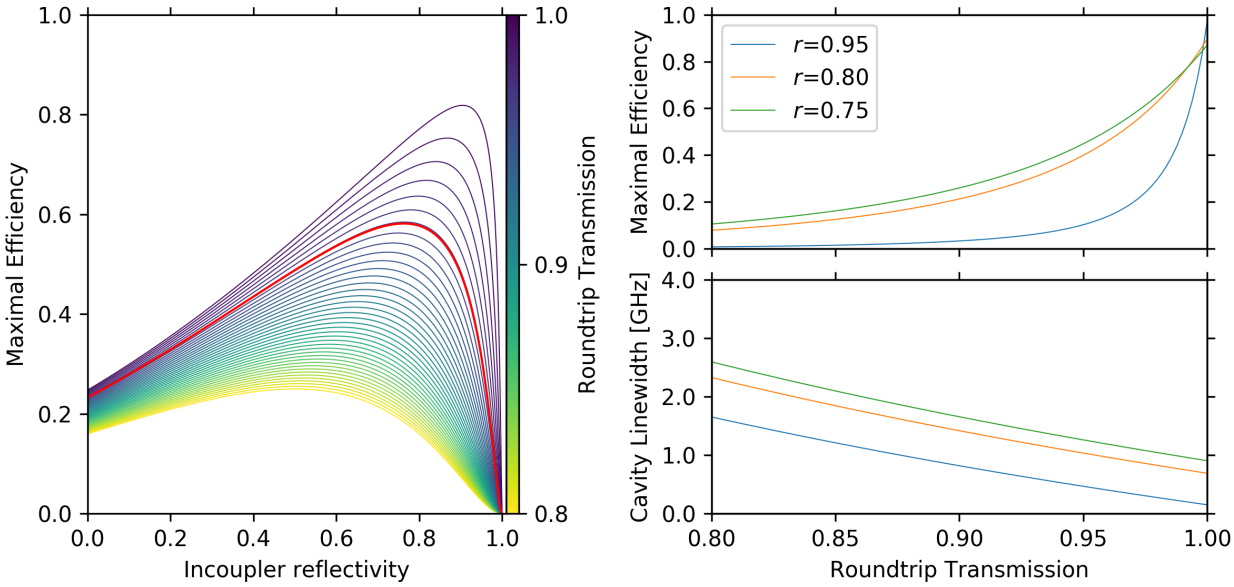


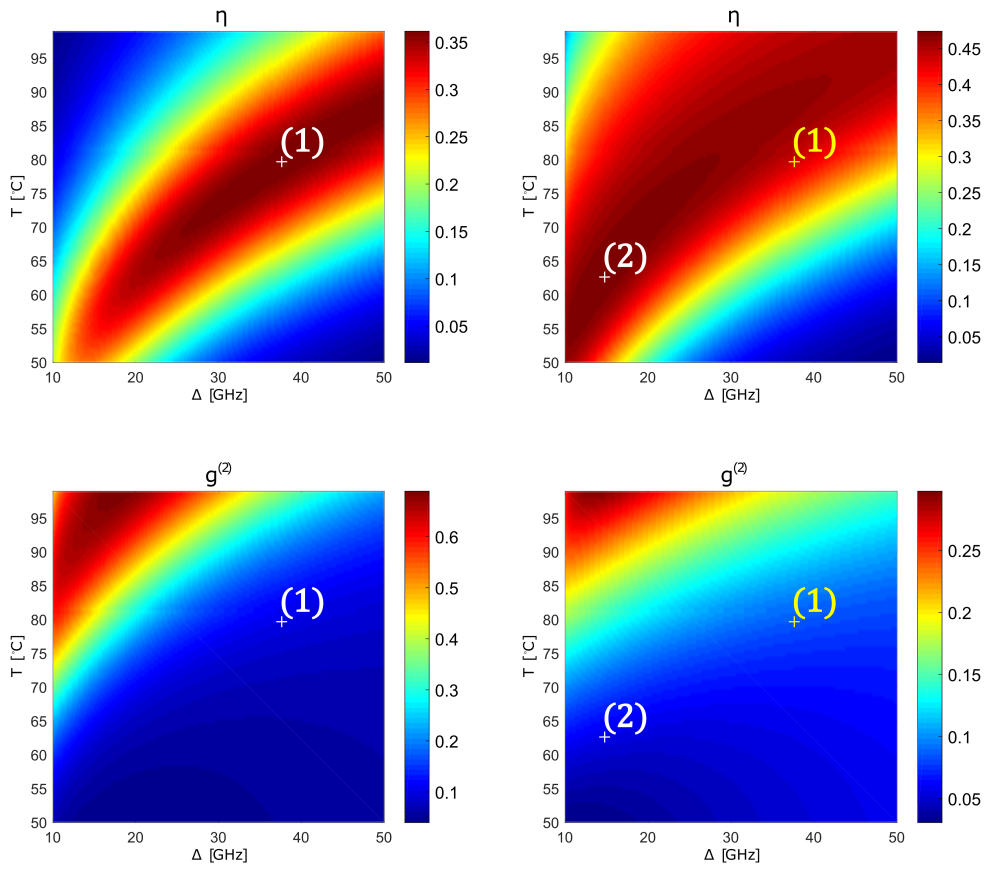
Figure 4.9: (Left) Asymptotic (i.e. loss limited) efficiency achievable as incoupler reflectivity is varied, for different values of intracavity loss (colour bar) – red line indicates the roundtrip transmission $\sqrt{.93}$ corresponding to the values in tables 4.5, and places an upper limit before the atoms are considered. (U.Right) Efficiency as intracavity loss is scanned for the three potential mirrors considered, (L.Right) Corresponding cavity linewidth for $M = 0$ cavity length. Note that the linear scale for the roundtrip transmission indicates that it becomes increasingly difficult access the highest efficiencies.

$g^{(2)}$ without constraining the intracavity loss for a conceivable next generation cavity configuration ($M = 2$, $r_{\text{in}} = 0.8$, $R = 3$m, with 10 Torr N_2). Here, we see that without optimal matching, to reach high efficiency we seek *high temperature* ($\sim 80°$C) and *large detuning* (38 GHz) enabling $\eta \sim 0.36$ and $g^{(2)} \sim 0.1$. If, however, we shape the control appropriately, we find conversely that it is more useful to operate at *lower temperatures* ($\sim 60°$C) and *closer to resonance* (15 GHz) and we can reach $\eta \sim 0.47$ and $g^{(2)} \sim 0.06$. For both of these cases, accounting for the pressure broadening and temperature – and the detuning, where we anticipate lifetimes of ~ 1 μs to be achievable. It is worth highlighting that the $\sim 50\%$ efficiency here is by no means a fundamental limit – this is obtained for a near-term design, using infrastructure that is readily available (for instance with similar losses) – much higher efficiencies may be achieved achieved with an optimised design.

The message to take from these results is that this is a high dimensional parameter space to search. In seeking to optimise on (i) efficiency; (ii) lifetime; (iii) noise suppression, along with (iv) bandwidth[15], we must consider the interplay of a range of parameters over which we have experimental control – to a greater or lesser degree – and the resulting trade-offs that arise. To highlight this, we distill the main parameters and their effects in table 4.6.

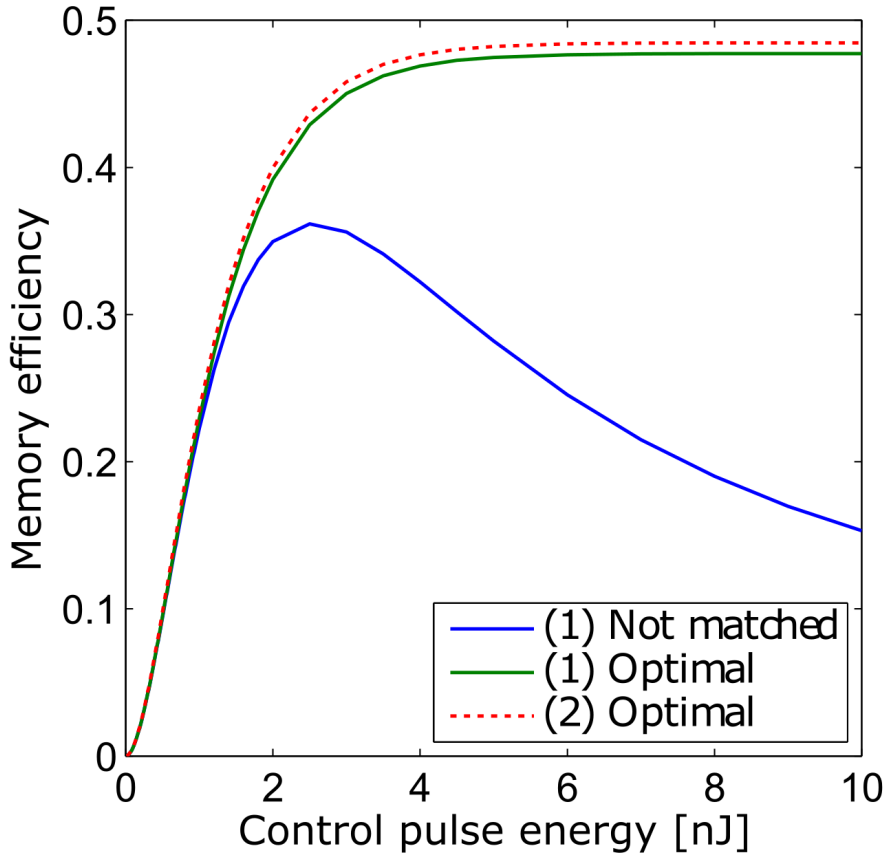
The methods presented here aim to facilitate the navigation of these parameters and their consequences. Using this one should consider the thresholds and targets for the four figures of merit highlighted – for the particular application in mind. This will then impose constraint on, for instance, what loss at the signal frequency is tolerable. In experimentally realising a given configuration, this model can then also be used to optimise on an experiment – given, say, the intrinsic loss of the intracavity components how we may best choose the temperature and detuning at which to operate. Although, this doesn't yet include the additional technical challenge of achieving simultaneous resonance for the control field.

[15]Where the bandwidth and suppression are related via the cavity finesse, via the free-spectral range for a given cavity.



(a) Not matched

(b) Optimal



(c) Efficiencies at the configurations highlighted in (a)

Figure 4.10: (a-b) Landscape showing efficiency & $g^{(2)}$ landscape for $M = 2$ cavity; the white crosses indicate the position of peak efficiency with and without optimal mode shaping of the control field. (c) variation of efficiency with control power: (1) corresponds to the configuration peak efficiency without mode matching (white cross left), (2) peak efficiency with optimally matched control (white cross right).

(a) Experimental Parameter:	(b) Physical Consequences:	(c) Figure of Merit:
Cavity order/length	linear absorption at signal	Efficiency
Incoupler reflectivity	cavity finesse	Lifetime
Backcoupler RoC	diffusion limited lifetime	Noise suppression
Intrinsic intracav. loss		Bandwidth
Cs temperature		
Laser detuning		
Rabi frequency		
Buffer gas (species & pressure)		

Table 4.6: Summary of (a) the control parameters which together give rise to (b) the main physical consequences, which in turn determine (c) the resulting figures of merit.

The general conclusion is that ultimately a low loss $M = 0$ cavity, with a large beam waist and sufficient buffer gas will be the best route for achieving simultaneously high efficiencies, long lifetimes and maximising noise suppression; whilst still allowing ∼GHz operation. However, reducing the losses sufficiently remains a technical challenge at present. These results also show that a realistic *near-term* implementation should reasonably be able to achieve efficiencies $\gtrsim 30\%$, with $g^{(2)} \lesssim 0.1$ with lifetimes $\gtrsim 1$ μs. Further improvements in efficiency are achievable if the intracavity losses can be reduced through sufficient engineering.

4.5 Summary

In this chapter I have outlined a series of investigations which seek to provide insights and tools for future iterations of the cavity Raman memory. We have identified a number of approaches for improving the cavity Raman memory, as well as highlighting several difficulties. In the final section here, I have sought to combine these elements and develop the tools to aid the navigation of these for future generations of the experiment.

4.5.0(a) Key Results

I have sought to address the means to optimise the performance of the cavity-Raman memory addressing three main themes: storage lifetime, memory efficiency and noise performance.

§4.2 Details of work that I have performed, including direct measurement of atomic diffusion in the presence of different buffer gases, and estimating the consequences on the memory lifetime that might be achievable. With this we consider the trade-off between memory lifetime and induced pressure broadening upon introducing different buffer gases.

§4.3 I briefly summarise parallel work led by S.E.T. investigating optical pumping – and the resulting memory efficiencies achievable; demonstrating enhancements in both by use of a molecular buffer gas enabling quenching.

§4.4 Presents an approach for performing a numerical parameter search, based upon simultaneously optimising on the efficiency and noise performance of the cavity memory for realistic components, including consideration of the mode matching between signal and control. This identifies the intracavity loss as a critical parameter we should seek to minimise.

CHAPTER 5

CONCLUSIONS PART I

In the first part of this thesis, I have presented an investigation into the feasibility of suppression of four–wave mixing in Λ systems, whilst enhancing the desired light–matter interaction by means of an optical cavity.

5.1 Summary

We have demonstrated the principle of operation of a noise suppressed Raman memory with single-photon-level coherent light. This showed a reduction of the noise present in the retrieved state as quantified by the μ_1 parameter. Whilst this does not access the quantum statistics of the retrieved state, this does at least indicate that the noise present is reduced with respect to the previous implementation.

We note that in addition to suppressed noise, introduction of the cavity offers enhancement available in terms of reduced overheads for instance in terms of the control field power needed, due to the cavity field-enhancement, as well as offering a reduced spatial footprint of the memory.

Based upon the theoretical description of the cavity Raman memory we have developed, I have numerically and experimentally investigated means of optimising the cavity Raman memory. This was considered in terms of three key themes: (i) memory lifetime; (ii) memory efficiency; (iii) reduced noise.

The first of these included an experimental investigation of diffusion with different buffer gases, and based upon these measurements projected the consequences for the memory lifetime[1]; and in conjunction with investigations of pumping efficiency by S.E.T., has guided our decision to move to N_2 buffer gas for future implementations of the Raman memory.

To address the second two of these themes, I applied the theoretical description of the cavity memory interaction developed by J.N. and myself, to numerically evaluate the performance of different potential *experimentally realistic* cavity configurations. As a consequence of these results, it is increasingly appealing to consider moving towards a machine-fabricated optical cavity, which as well as offering improved stability, and reduced loss, would constitute a readily integrated and scalable device – additionally requiring reduced experimental overheads – and size – with respect to the bulk

[1] If operating with more narrowband pulses, one could employ an anti-relaxation coating, such as paraffin, which has been commonplace in precision magnetometers (e.g. [177]), however for this to be beneficial one requires long pulses such that the light can interact with all atoms in the cell over extended time periods as they diffuse into and out of the beam, without suffering dephasing collisions with the cell wall.

optics implementation.[2]

At the same time, bearing in mind the intracavity loss as a limiting factor, I considered the interplay between accessible control parameters – namely the detuning, cell temperature and control field power – and given a practical cavity configuration investigated sought to optimise upon key metrics of the memory performance. The findings from this suggest that, without optimal mode matching between the signal and control, it could be favourable to operate at increased detuning, and compensate the reduced coupling with an increase in atom number; however, if one is able to appropriately shape the control for a target signal shape *at a given control pulse energy* it is instead beneficial to operate closer to resonance at lower temperatures.

Although this proof of principle demonstration is not yet sufficient for single photon inputs, firstly for the levels of suppressions achieved: a $\mu_1 \sim 0.2$, means that a single photon from a downconversion source with a heralding efficiency of 0.25 would only just surpass the noise; a precision monolithic microcavity would enable much greater suppression factors to be reached. In addition to this facilitating higher efficiencies to be reached, through reduced intracavity loss, as well as improving the stability of the experiment over the long durations needed for acquisition of photon statistics.

The theoretical description of the system has also identified the single-modeness of the memory interaction. In contrast to multi–mode capacity, which is desireable for some memory applications [64], we turn to single mode storage in the next part of the thesis.

5.2 Wider Landscape

In this thesis so far I have introduced the cavity Raman memory. On the basis of our experimental proof-of-principle demonstration, together with the optimisation outlined in the previous chapter, this platform has the potential to be broadband; low noise; high efficiency; and could constitute a readily scalable system with a small physical footprint. In fact with the $M = 0$ cavity, considered as a promising configuration, a memory on this scale could concievably be produced, with a ultra-thin vapour cell (e.g. as might be used for atomic vapour based sensors [178]), into a chip-scale integrated package. In this case one might imagine incorporating the optics into the cell so as to minimise the number of interfaces. Additionally, with lower operating temperatures possible with optimal pulse shaping of the control, spin preserving coatings, such as paraffin, may be employed, since these are limited to temperatures below $\sim 60\,^{\circ}\mathrm{C}$ [179].

These results provide another approach to optimising the performance of a quantum memory, without moving towards "cold systems", whether this is in the solid state, or cold atoms, where typically noise suppression and efficiency are enhanced, at the cost of increased experimental overheads and reduced bandwidth.

With regards local synchronisation however, there are alternative, perhaps more promising, protocols for memories intended for local (i.e. short time-scale) synchronisation applications, such as the ladder protocols in warm vapour ORCA [115] and FLAME [116], or the diamond Raman memory [131]. These are typically extremely low noise, and broad-band; although at present limited in efficiency.

Employing computational methods to optimise experimental parameters is an appealing approach. The use of a deep-learning neural-network has been employed by a group at the ANU in order to

[2] Although, already the warm alkali vapour Raman memory is a relatively compact set up, without requiring cryogenics or vacuum systems.

optimise the experimental parameters for the loading of a magneto-optical trap [180]; and indeed if an appropriate cost function can be defined in terms of the tolerable noise level and target efficiency, such approaches could be used here.

Before one can get to that stage, certain parameters must be constrained – say in our case the buffer gas pressure, and the numerical (and experimental) investigations here are the first step.

One might consider additional applications based upon these methods of enhancing desireable interaction, whilst suppressing others; for instance developing upon studies of the three-mode light-matter interactions in alkali vapours between the Stokes, anti-Stokes and spin wave [181]. Incorporating this with our results one could imagine a platform which might enable independently manipulating the "beam splitter" (quantum memory) and "two-mode squeezing" [182] parts of the interaction, and ask if this could offer enhancements for networking applications, and which figures of merit are most relevant.

5.3 Moving Forward

In this thesis thus far we have developed the tools for the design of systems to manipulate the light-matter interaction by means of an optical cavity; and demonstrated this in practice by showing the suppression of noise in the Raman memory. The problem/solution however, is a general one.

5.3.1 Improved Cavity Design

Moving forwards experimentally from the results of the previous two chapters, which was based upon an $M = 2$ design, such that it could be constructed with bulk table top optics, the next step is to operate a single port cavity in reflection, such that the achievable efficiency is not degraded by read-out from multiple ports. For near-term designs standard high reflectivity optics can be used for the two flat mirrors (near 45° incidence), and all other optics are near normal incidence, and in-couple through the back of the curved mirror. This eliminates intracavity birefringence from the cavity optics. With readily available components for a number of potential designs with which, in conjunction with the numerical package presented, we have the tools to experimentally explore the parameter space such that we can optimise on memory efficiency, lifetime, and noise suppression.

5.3.2 Alternative Approaches to Noise Suppression: "Magic Λ"

Suppressing FWM with a cavity, is not the only way that one can modulate the density of states. A more recent approach under investigation in the memories lab in Oxford is based upon the principle of introducing an absorptive feature seen by the mode to be suppressed. Whilst this approach has been demonstrated in rubidium by introducing a different isotope [183] which has an atomic transition resonant with one of the optical fields yielding a suppression of the FWM interaction; in caesium there are no other stable isotopes that could be admixed. However, by choosing the detuning and pumping configuration appropriately we can make it such that the unseeded arm, or parasitic mode, of the FWM interaction is resonant with the populated atomic transition, as shown in figure 5.1. This constrains the detuning of the memory process to be equal to the ground-state splitting.

Indeed in the future, multiple methods could be employed, in parallel with the cavity memory with the view to enhancing the memory performance further.

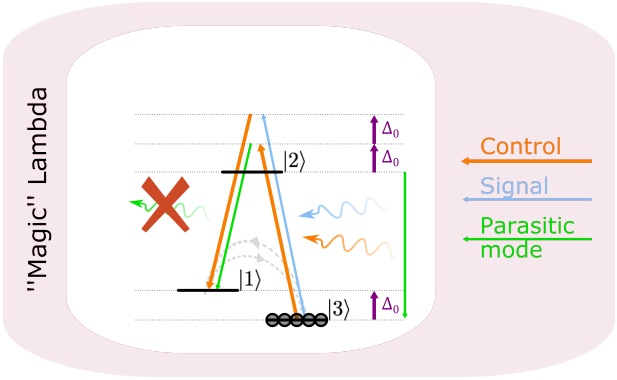


Figure 5.1: Pumping and detuning configuration such that the unwanted mode of the FWM is resonant with the populated atomic transition, and is subject to a highly absorptive feature.

5.3.3 Temporal Modes

Temporally shaping the control pulse as a means to optimise the memory performance has been repeatedly highlighted as an enabling feature to help the Raman memory become a more useful device. This broadly follows two arguments: that of enhancement of light-matter interaction, and secondly exploiting the single mode nature of the memory interaction.

The first of these is a general feature of any light-matter system, and involves tailoring the properties – in particular in temporal amplitude and phase – of the optical fields to enhance the interaction strength with a matter system. This could be adopting the optimal shape for a photon to be absorbed by an atom or atom-like system (i.e. a rising exponential – by arguments of time-reversal [139]), or as in our case in which the control field mediates the interaction between a separate optical field and the matter system. And, as we have shown, this is necessary for getting the most out of the cavity Raman memory.

The latter is an interesting feature, since often the multi-mode capacity of a memory is focussed on as a desirable feature [64, 111] but in the next part of this thesis I explore the applications exploiting the mode selectivity of the Raman memory for the manipulation of temporal-modes of light.

Moving to a carved control pulse carved from a diode laser addresses both of these, whilst also removing the experiment's dependence on the Ti:Sa, and provides another means to reduce the experimental overheads.

We now turn to an investigation of the applications of the Raman memory for interfacing with temporal modes of light.

PART II

TEMPORAL MODES OF LIGHT

CHAPTER 6

TEMPORAL MODES OF LIGHT: AN INTRODUCTION

6.1 Making use of Temporal Shapes in Light

As discussed in the introduction to this thesis: practical, *scalable*, architectures for quantum information processing and quantum communication will at some point rely upon hybrid systems drawing upon the strengths of different media: for instance in which one exploits the precision and controllability of quantum states and interactions in material systems and photonic systems which exhibit only weak interaction with their environment (and other photons). The former facilitate information processing, whereas the latter are ideal for the transmission of information between nodes of a quantum network.

Interfacing optical states with material systems have requirements on the spectral/temporal properties of the light. Therefore, in order to build up a larger scale quantum network or computation by concatenating processes in these different media, a key requirement is the ability to coherently transform the photon in time and frequency in order to interface different systems.

6.1.1 Temporal Modes for QIP

Temporal modes (TMs) of single photon states have recieved increasing interest recently as an appealing platform for quantum information processing [184–186].

To highlight some of the appeal of TMs as a platform for quantum information in photonic states, we first consider the degrees of freedom of a photon in which information may be encoded: (i) helicity; and (ii–iv) three components of momentum vector. In a beam geometry these can be interpreted as polarisation, transverse profile (2 dimensions) and lastly frequency. Polarisation is the most easily accessible and traditionally exploited degree of freedom but this is limited to a two dimensional Hilbert space, in which information is encoded in qubits (two state systems). Transverse profile has received much attention in recent years with orbital angular momentum (OAM) states [187, 188] offering a high-dimensional set of basis states allowing one to work instead with information encoded in qu*d*its, in a d dimensional Hilbert space.

Frequency (and its Fourier-conjugate, time) is an underexploited degree of freedom. Whilst discrete time/frequency–bin encoding (e.g. [189]) offers a simple basis to interact with, these quasi-orthogonal states do not make optimal use of time-frequency space. Consequently the dimensionality available is fundamentally limited by the time/frequency window available for the encoding of one qudit since one must use temporally short or spectrally narrow pulses to define a series of distinct,

quasi-orthogonal channels.

In contrast to this, one may define wave-packet modes which overlap in time and frequency but are *field orthogonal*, referred to as "temporal modes" (TMs). TMs form a complete basis in which an arbitrary state, in the energy degree of freedom, can be decomposed. Analogously to OAM states, TMs overlap in time/frequency; however, since they occupy the same spatial distribution they are naturally compatible with waveguide devices and single mode fibre (SMF) networks. Furthermore, due to their overlapping spectra, TMs are insensitive to static or slow-varying perturbations such as dispersion which identically affects all modes. In principle, these TMs suffer no cross talk between different modes, whilst maintaining optimal information packing density in some region of time/frequency space.

Temporal mode encoding can be applied to any quantum state of light, one such choice is single photons [184], although one could equally consider, for instance, quadrature squeezed states[1] [190].

6.1.2 Implementing TM encoding

We choose Hermite–Gaussian (HG) Polynomials as our complete set of basis states onto which any state may be decomposed, as the HG polynomials closely resemble the eigenmodes of photon pair sources [191].

To describe these states we can write down the state of a photon in a particular temporal mode as a coherent continuum superposition of single photon states in different monochromatic modes or equivalently at different creation times, where the two are related by a Fourier transform.

In the time domain, the TM corresponding to a temporal wavepacket $f_j(t)$, describing the TM j, could be written:

$$\begin{aligned} |A_j\rangle &= \int \mathrm{d}t\, f_j(t)\hat{A}^\dagger(t)\,|0\rangle \\ &\equiv \hat{A}_j^\dagger\,|0\rangle\,, \end{aligned} \tag{6.1}$$

where $\hat{A}(t)$ is the annihilation operator for a photon at time t, and $\hat{A}_j$ is a collective mode operator annihilating the wave packet state $|A_j\rangle$.

As these are a complete basis set, an arbitrary state $|\Psi\rangle$ may be decomposed as a superposition of these basis states $|A_j\rangle$

$$|\Psi\rangle = \sum_{j=0}^{\infty} c_j \hat{A}_j^\dagger\,|0\rangle\,. \tag{6.2}$$

These overlapping states are then orthogonal with respect to the temporal (frequency) integral, where the inner product on this space is

$$\begin{aligned} \langle A_j|A_k\rangle &= \iint_{-\infty}^{\infty} \mathrm{d}t'\mathrm{d}t''\, f_j^*(t') f_k(t'')\,\langle 0|\hat{A}(t')\hat{A}^\dagger(t'')|0\rangle \\ &= \int_{-\infty}^{\infty} \mathrm{d}t'\, f_j^*(t') f_k(t') \qquad \left(\,\langle 0|\hat{A}(t')\hat{A}^\dagger(t'')|0\rangle \;=\; \delta\left(t'-t''\right)\,\right) \\ &= \delta_{j,k}, \end{aligned} \tag{6.3}$$

by the orthonormality of the chosen TM basis states. And finally, by construction, these TM ladder

[1] Whilst these states would also be compatible with the quantum memory considered here, further investigation would be required to verify the tolerance of such states to finite memory efficiency and other inherent losses.

operators satisfy Bosonic commutation relations:

$$\begin{aligned}\left[\hat{A}_j,\, \hat{A}_k^\dagger\right] &= \iint \mathrm{d}t\,\mathrm{d}t'\, f_j^*(t) f_k(t') \left[\hat{A}(t),\, \hat{A}^\dagger(t')\right] \\ &= \iint \mathrm{d}t\,\mathrm{d}t'\, f_j^*(t) f_k(t') \delta(t-t') \\ &= \delta_{j,k}. \end{aligned} \tag{6.4}$$

Thus the $\hat{A}_j$ can be seen to correspond to a broadband-mode operator.

We must next consider how one interacts with these states. To illustrate this, one may compare with other encodings:

Encoding basis	Addressing	Manipulation
Polarisation	polarisation-sensitive beam splitters (PBS)	Waveplates
Frequency (DWDM)	frequency-sensitive beam splitters	Frequency conversion
Spatial (OAM)	spatially varying reflectivity (e.g. SLMs)	Spiral phase-plate

so by analogy, the manipulation of temporal modes necessitates being able to implement time non-stationary beam splitters: that is, one requires time non-stationary elements [192], since the information is encoded in temporally varying waveforms.

6.1.3 Enhanced Light–Matter Interaction

Another motivation for working with temporal modes, where we have control over the temporal amplitude and phase of an optical signal, is to use these shapes to optimise interaction between the light and some material system.

We mentioned the stringent requirements on the photon properties for interaction with various matter systems; one instance of this which we have already touched upon is the optimal profile of the control field in the Raman memory interaction. As the Rabi frequency is increased, the efficiency becomes limited by the a.c. Stark shift – where the increased control field power induces a shift in the atomic energy levels and drives the system off two-photon resonance. With control over the temporal shape of the control field amplitude and phase, this limit may be surpassed by dynamically modulating the control field in time such that the two-photon resonance condition is retained.

To illustrate this, we take an example for the cavity Raman memory. In figure 6.1 I show the optimal control field temporal profile which is needed to address the first three HG basis states, as the pulse energy ($\propto W = \int_\infty^\infty \mathrm{d}t' |\Omega(t')|^2$, where $\Omega(t)$ is the Rabi frequency) is varied. At lower powers the optimal control field closely follows the target signal to be addressed; as W is increased, the distortion in temporal amplitude and phase that arises as the Rabi frequency initially begins to increase (at the rising edge of the pulse) enables the detrimental effects of the a.c. Stark shift to be counteracted. Recalling that the cavity memory is a single mode interaction, these control profiles are those for addressing *solely* the target signal temporal modes, with efficiencies determined by the cavity parameters. These are derived from inverting equation 3.111.

6.2 Devices for Temporal Modes

To interface with temporal modes, and selectively address these overlapping TMs, the necessary dynamics may be interpreted as that of a time non-stationary beam splitter (as argued above). In

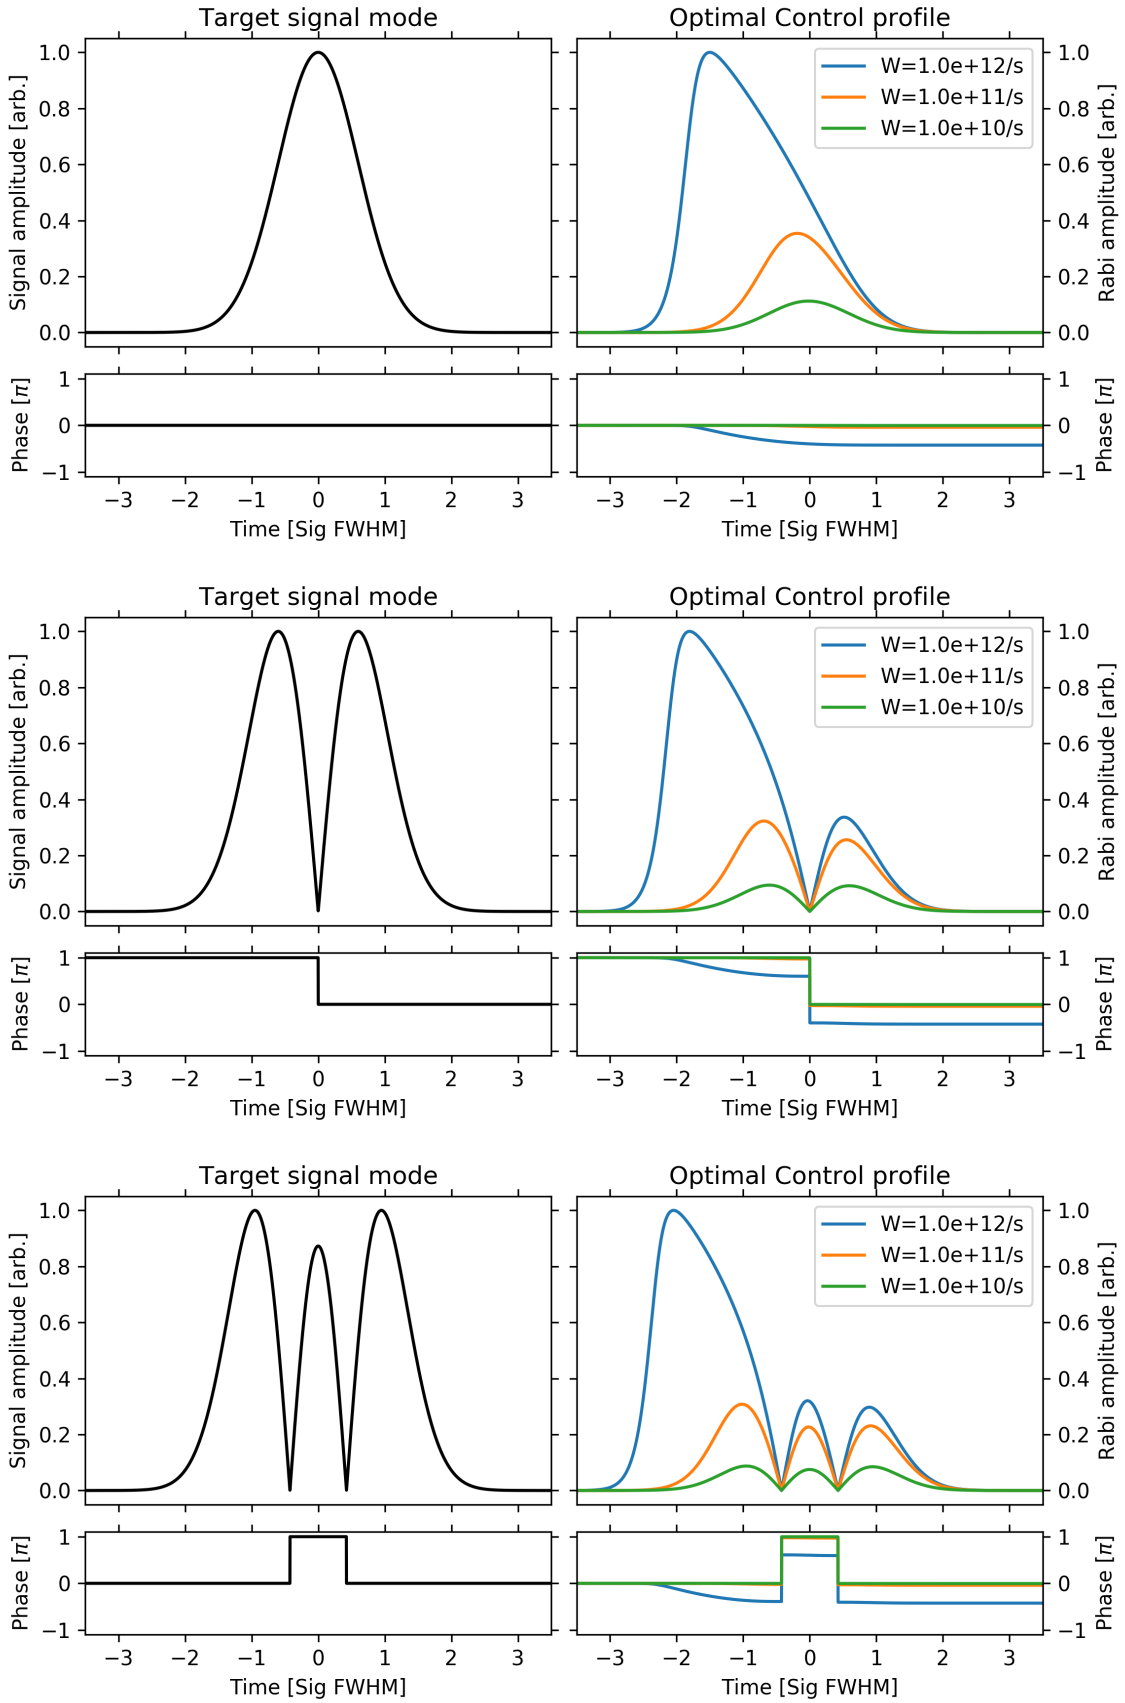


Figure 6.1: Ideal control field temporal profiles for the storage of target signals for the first three Hermite-Gaussian basis states interfaced with the cavity Raman memory, for the cavity configuration shown in figure 4.10 at the non-mode matched optimum ($T = 80°$ C, $\Delta = 38$ GHz).

general this can be realised by having some time dependent coupling mediated by some external control field.

If one considers an incident qudit state, decomposed in the TM basis:

$$|\Psi_{\text{in}}\rangle = \sum_{j=0}^{d-1} c_j \, |A_j\rangle \, . \tag{6.5}$$

If this is acted upon by a TM beam-splitter, which with appropriate tuning, could be configured to address some state $|A_n\rangle$ and map this into some other mode B_0. This then enacts:

$$\mathcal{B}_0 \sim \theta_n \left(\hat{A}_n^\dagger \hat{B}_0 + \hat{A}_n \hat{B}_0^\dagger \right), \tag{6.6}$$

with an efficiency parametrised by a "reflectivity" $r_n = \sin(\theta)$. Here B_0 represents some "register mode" into which a chosen mode A_n is mapped.

This is depicted schematically in 6.3 whereby one can envision an input state encoded in a superposition of TMs, which may be coherently mapped into a set of different modes; analogously to a polarisation superposition being mapped by a PBS into two spatial modes.

6.2.1 Quantum Pulse Gates

It was first proposed by Raymer in 2010 that sum-frequency generation (SFG) could be exploited as a "two-colour beamsplitter" [193]. This idea was then developed by Silberhorn's group in Paderborn to use group-velocity matched SFG as a "Quantum Pulse Gate" (QPG)[194, 195], and since then the QPG has been adopted by other groups [196].

In this scheme a signal is sent into a non-linear medium, which is simultaneously pumped by a control pulse, in this process the control field *selectively* frequency converts a part of the signal field, determined by the control profile. The $\chi^{(2)}$ non-linearity results in a frequency shift and reshaping of the separated mode, governed by the phase matching conditions. See figure 6.2 for a cartoon of the operation schematic of the pulse gate.

In practice, the QPG relies upon group-velocity matching between the signal and pump which is not always straightforward to realise, due to the dispersion properties of non-linear crystals. Furthermore due to phase-matching conditions, there are also limitations in spectral bandwidth on which a QPG may operate, which results in a lower limit in the region ~ 100 GHz [197].

For compatibility of the TM framework with solid state or atomic systems, one is limited to signal bandwidths in the MHz-GHz regime. The Raman quantum memory can provide this functionality.

6.2.2 Single Mode Quantum Memories

The analogy of the Raman memory as a light–matter beam splitter has been long understood [106, 138] (refer back to equation[2] 2.47) however, the application of the time dependent beam-splitter for interfacing with temporal modes is relatively recent.

The feature of temporal mode selectivity in this system was highlighted in the derivation of the Raman memory equations of motion when investigating the suppression of FWM by means of a cavity

[2] Although, as shown, this is only part of the full interaction and (particularly at high coupling strengths) other effects may be present.

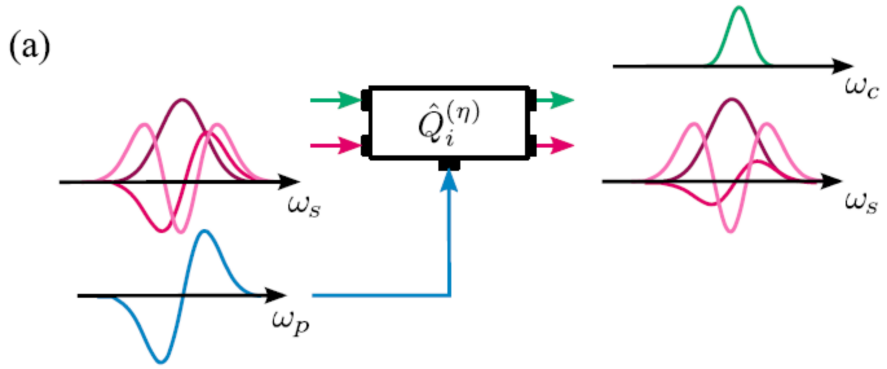


Figure 6.2: Schematic of quantum pulse-gate (QPG) operation, repeated (with permission) from [184].

[147] (or see equation 3.55); where the cavity enables a weak coupling on a single pass to be enhanced over many round trips.

The framework of temporal modes as a platform for QIP shows that having the ability to selectively address particular temporal modes of a state is an elementary building block. Indeed, this means that we may exploit single–modeness of our memory to our advantage, in contrast with focussing on multi-mode capacity of a quantum memory, as has been a goal of the wider memories community (e.g. for repeater applications [64, 111, 126] or the storage of images [198, 199]).

In contrast to the QPG the Raman memory enables operations with TMs in a spectral bandwidth regime not accessible in SFG due to phase matching conditions. Furthermore, with a memory, however, our interaction is not limited to separating out orthogonal TMs, but also enables their manipulation. Once information encoded in TMs has selectively been addressed and written into the storage medium, the QM also constitutes a system in which the signal can be retrieved into a different mode and hence this may be used to interface between different encodings.

To clarify this: with full control of the read-in and read-out control pulses, we are able to select one mode from an input, and retrieve it into another mode. In the simplest case one could envision using TM encoding (exploiting the high dimensionality available within a single spatial mode, for transmission down fibre, for example) and using a quantum memory to selectively delay different components – thereby mapping from TMs to a time–bin encoding. Aside from implementing a delay of a target temporal mode, the memory also facilitates more involved manipulation of the complex spectrum of the pulse. This has been demonstrated for spectral shearing [119]; however more developed manipulation still remains to be demonstrated. In part this is due to the difficulties in both the generation and characterisation of the pulses mediating the interaction for such operations.

Furthermore, using the quantum memory for the purposes of manipulating TMs is also therefore appealing as it simultaneously enables the conversion between a range of different encodings. We compare this to implementations with the QPG which to achieve reshaping requires the concatenation of *multiple devices*[3] [195]. Additionally, using a quantum memory for these applications also enables other features, such as the "user specified delay" of the retrieved signal, which as well as enabling a straightforward mapping to time bin encoding described above, could offer additional functionality allowing for electronic feed forward and switching downstream, so that different operations may be

[3] Although, with the memory one requires *two interactions* (i.e. read-in and then read-out), albeit in a single physical device.

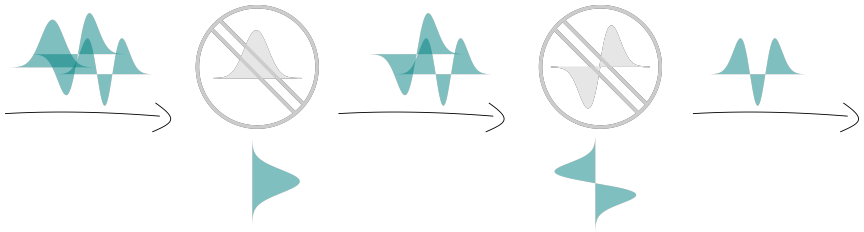


Figure 6.3: Cartoon of TM selective operations

applied to each component of the qudit.

6.3 Outlook

In the following chapter we perform a preliminary experimental and theoretical characterisation of the modal structure of the memory interaction, and seek to generate new insights into the application of the Raman quantum memory as an interface for the manipulation of temporal modes of light.

To achieve this, we contribute to the toolbox of methods which facilitate interaction with temporal modes – from the generation and verification of the states, selectively addressing particular temporal modes within them, and investigating the means to characterise the performance via tomographically reconstructing the process enacted by the device.

This then provides the basis for future work to use the Raman memory as a flexible device enabling the wider applicabilty of TMs within the landscape of platforms and schemes for quantum information processing.

CHAPTER 7

MODE–SELECTIVITY IN A RAMAN MEMORY

In this chapter I present the initial findings of our group's investigations into TM selectivity mediated by a Raman memory.

The project has been overseen by B.B., whilst I built, maintained and operated the pulse carving and characterisation infrastructure. For interfacing these shaped pulses with the memory, S.E.T. maintained and operated the caesium memory, and the detection part of the experiment. We were both present equally for the acquisition of memory data. The analysis of these data have all been conducted by myself, with some advice from J.N. and B.B., where I have developed existing tomographic techniques from a range of references, cited where appropriate, adapting them to our experiment.

7.1 Introduction

On the basis of our theoretical model of the cavity memory outlined in chapter 3, we draw the conclusion that the cavity Raman quantum memory constitutes a system which enables mode-selective operation, whereby we can selectively interact with a single temporal–mode of our choosing, whilst leaving all orthogonal components unaffected.

The first steps to realising the Raman memory as a useful device for the manipulation of TMs is to investigate the storage of different input signal temporal modes for a *given control field* profile. This step requires that we are able to prepare arbitrary input states, and that we can verify those states produced (see section 7.3.1 and appendix B); and then measure the efficiency with which they are written into the memory.

Once this is achieved, we then must show that the interaction itself is single mode – and that we do indeed address just one temporal mode. To do this we employ tomographic techniques based upon our knowledge of the system and the measured read-in efficiencies for each state in order to reconstruct the Green's function of the memory interaction (see 7.4.2). Having obtained the Green's function, we can then extract the mode structure of the "atoms plus control field" (A$-\Omega$) system: that is, the eigenmodes of the interaction, and the efficiencies with which they are addressed. A truly mode selective device will have only one non-zero singular value.

To experimentally demonstrate this, we start with an implementation of the Raman memory in the bulk setup (sans cavity) for experimental simplicity, to probe the resulting behaviour with coherent states[1] prepared in a range of temporal shapes.

[1] N.B. here, we are interested in the *modal properties* of the interaction, we do not require non-classical input states.

To reconstruct the Green's function of the memory interaction, i.e. the A$-\Omega$ system, we consider the efficiency with which different states are written into the memory for a fixed control field profile. By recording the read–in efficiency for different, known, probe states we may map out the Green's function of the memory process.

In this chapter, therefore, in addition to demonstrating the first steps of realising the Raman memory as a system for the manipulation of temporal modes of light, we develop the toolbox of techniques for the characterisation and verification of the structure of the interaction.

7.2 Concepts

7.2.1 Tomographic Analysis

To informationally probe a quantum system, one must have a set of states which span the domain containing the system – that is, a basis onto which the system can be decomposed without any loss of information. The optimal choice here is to construct mutually unbiased bases (MUBs), where these are choices of bases such that if a state is encoded in one basis, but measured in another the resulting outcome is *unbiased* between the measurement outcomes in that basis, i.e. yielding a probability of $1/d$ for any outcome when measured in the incorrect basis, where d is the dimension of the Hilbert space $d = \dim(\mathcal{H})$.

To illustrate this in the simple case of a qubit, measurements in only one basis, say Z, will only yield information about the populations. One needs measurement in a complementary basis in order to access information about the off-diagonal elements.

Let us denote our basis states as $|\phi^{(T)}_{\alpha,\beta}(\tau)\rangle$ where α denotes the choice of basis, β the dimension index, and τ the relative time delay between the signal and control pulse arrival time at the memory. The superscript T is included to denote the time of the measurement in the lab. For instance, the familiar case of a two dimensional Hilbert space for a qubit we may choose three MUBs, for instance (but not uniquely) using the construction in table 7.1, but this can be extended to arbitrary dimension, where in general for a d-dimensional Hilbert space, there are (at most) $d+1$ MUBs [200].

α		0	1	2
β	0	$\|0\rangle$	$\|+\rangle = \frac{1}{\sqrt{2}}(\|0\rangle + \|1\rangle)$	$\|\mathrm{i}\rangle = \frac{1}{\sqrt{2}}(\|0\rangle + \mathrm{i}\,\|1\rangle)$
	1	$\|1\rangle$	$\|-\rangle = \frac{1}{\sqrt{2}}(\|0\rangle - \|1\rangle)$	$\|\mathrm{i}^*\rangle = \frac{1}{\sqrt{2}}(\|0\rangle - \mathrm{i}\,\|1\rangle)$

Table 7.1: Construction of the three mutually unbiased bases in a two dimensional Hilbert space.

With (i) many identically prepared copies of the unknown quantum system, and (ii) a series of measurements constructed from these mutually-unbiased states, one has access to sufficient information to enable the complete reconstruction of a state. In general this requires measurement of at least $d+1$ non-commuting observables; hence observables corresponding to MUBs are optimal as this minimises the number of non-commuting observables to measure [201].

We may exactly apply this to the temporal–mode basis as in section 6.1.2, in which our qudit encoding represents the value n in the n^{th} order Hermite–Gaussian (HG) function. These modes constitute field orthogonal states. In the following sections we refer to this as the computational (COMP) basis, and the d mutually unbiased bases (MUBj, for $j = 0, \ldots, d-1$).

7.2.2 Modelling the Memory Interaction

To theoretically complement the experiment, we can simulate the predicted memory efficiency given estimates of our experimental parameters. In the present work we use the simplest model of the Raman memory – that is "ideal memory" operation, sans FWM[2] – as used in [134] to characterise the system:[3]

$$\begin{aligned} B(z) &= \int_{-\infty}^{\infty} \mathrm{d}t\; K(z,t) S_{\mathrm{in}}(t) \\ K(z,t) &= \mathrm{e}^{\omega(t)W/\Gamma} J_0\left[2C\sqrt{\omega(t)(1-z)}\right], \end{aligned} \tag{7.1}$$

where the read–in efficiency is taken to be $\int \mathrm{d}z\; B^\dagger B / \int \mathrm{d}t\; S_{\mathrm{in}}^\dagger S_{\mathrm{in}}$. The complex pre–factor is responsible for the Stark–shift, and in the far detuned (Raman) limit, $\Gamma \approx -\mathrm{i}\Delta$, which otherwise may also give rise to linear absorption of the signal field.

We can define a corresponding transformation, within this model [133], for the mapping between the incident and transmitted optical modes (in the first time bin) – i.e. to obtain the optical field transmitted:

$$\begin{aligned} S_{\mathrm{tr}}(t) &= \int_{-\infty}^{\infty} \mathrm{d}t'\; L(t,t') S_{\mathrm{in}}(t') \\ L(t,t') &= \delta(t-t') - C\Theta(t-t') \frac{J_1\left[2C\sqrt{\omega(t)-\omega(t')}\right]}{\sqrt{\omega(t)-\omega(t')}}, \end{aligned} \tag{7.2}$$

where with negligible absorption (and negligible FWM gain), we must have that $K^\dagger K + L^\dagger L = \mathbb{1}$ i.e. that the photon is either mapped to the atomic excitation or transmitted.

7.2.2(a) Numerical Methods

To numerically describe the memory interaction, given a known temporal mode of the input signal, we must perform numerical integrations of the form in equation 7.1.

We take a brief diversion to describe the Nystrom method for numerical integration [see e.g. 202, chapt. 19], whereby we approximate the integral with a weighted sum:

$$\int_a^b \mathrm{d}s\; y(s) = \sum_{j=0}^{N-1} w_j\, y(s_j), \tag{7.3}$$

where at the N discrete points on the integration domain $[a,b]$, the values $\{w_j\}$ are integration weights, and the integrand is sampled at the abcissas $\{s_j\}$.

Quadrature methods for numerical integration can be summarised as *adding up the value of the integrand at a sequence of abcissas, where we seek to obtain the integral as accurately as possible, whilst minimising the number of evaluations of the integrand.* To do this, the method is based upon mimicking our integrand as a weighted sum of mutually orthogonal polynomials. For instance, a suitable choice [see e.g. 203, chapt. 27] is the Legendre polynomials $P_l(x)$ on the domain $x \in [-1,1]$.

With this construction, matrix operations – or integrals – are modified to include the quadrature

[2] In the limit, also, that $\Delta \gg \delta_c,\, \Omega\, \gamma$, where Δ is the detuning, δ_c – bandwidth, Ω – Rabi frequency, and γ – linewidth.
[3] Summarised earlier in §2.3.

weights. Transforming the elements of matrices and vectors according to:

$$K(z_i, t_j) \to K' = \sqrt{w_z^{[i]}} K(z_i, t_j) \sqrt{w_t^{[j]}}$$
$$S(t_j) \to S' = \sqrt{w_t^{[j]}} S(t_j),$$

such that

$$\left[\boldsymbol{K'S'}\right]_i = \sqrt{w_z^{[i]}} K(z_i, t_j) w_t^{[j]} S(t_j) = [\boldsymbol{B'}]_i, \tag{7.4}$$

and the resulting vector is exactly the quadrature approximation for the required integral.

We therefore perform all calculations in this transformed space (and the primes $'$ are assumed). Note that, when extracting the eigensystem of a matrix, the eigenvector must be transformed back by applying the inverse of the weights.

The following analysis is peformed using Gaussian quadrature with abcissas and weights derived from the Legendre polynomials. Some discussion of the choice of sampling is given in appendix E.

7.2.2(b) Modelling Memory Performance

With this model in place we can simulate the projected performance of the memory under a given set of experimental conditions (i.e. for some A$-\Omega$), and in particular ask the following questions:

- what is the memory efficiency given our prepared signal mode?
- what is the *maximal* memory efficiency, if we were to prepare a signal in the optimal mode?
- what is the single–modeness of the interaction in this regime?

We consider the structure of K in terms of two independent components (i) the number of atoms[4] $\sim d\gamma$; and (ii) the control field (shape ω, and magnitude W).

In practice, W is the most experimentally accessible parameter, by simply varying the control field power incident on the memory by means of a waveplate. In principle the atom number is also accessible; however, the additional time overheads for varying the cell temperature allowing for equilibration and measurement of the optical depth mean that here we focus on the Rabi frequency.

Performing a singular-value decomposition (SVD) on $\boldsymbol{K}$, or symmetric eigenvalue decomposition on $\boldsymbol{K}^\dagger\boldsymbol{K}$, we may extract the eigenmodes of the memory interaction. By construction $\boldsymbol{K}^\dagger\boldsymbol{K}$ is Hermitian, and hence yields real eigenvalues and orthogonal eigenvectors. The eigenvalues here correspond exactly to the memory efficiency, i.e. the read–in efficiency may be evaluated as

$$\begin{aligned} N_B &= \left\langle \boldsymbol{B}^\dagger \boldsymbol{B} \right\rangle \\ &= \left\langle \boldsymbol{S}_{\text{in}}^\dagger \boldsymbol{K}^\dagger \boldsymbol{K} \boldsymbol{S}_{\text{in}} \right\rangle, \end{aligned}$$

thus, if $\boldsymbol{S}_{\text{in}}$ is shaped such that it corresponds to an eigenmode of the memory interaction with eigenvalue λ_i; using the matrix properties above, the storage efficiency comes out as $\lambda_i^2 = \eta_i$.

We first consider the separability [196] as a metric of the single modeness of the interaction where:

$$\mathbb{s} = \lambda_0 / \sum \lambda_k, \tag{7.5}$$

[4] Confusingly, perhaps, the quantity $d\gamma$ is *independent* of the linewidth γ. See equations 2.54 and 2.34.

and λ_0 is the largest eigenvalue. Another metric is the "selectivity" $S = \lambda_0 \, \mathbb{s}$ which encompasses both the degree of separability of the Green's function and the overall process efficiency [196].

Note that complications arise in the presence of FWM gain and efficiencies are no longer bounded by unity, as discussed in section 4.3; however, the simplified model used here in equation 7.1 does not include the FWM interaction. We therefore restrict the following analysis to a weak interaction regime where the beam-splitter part of the interaction is dominant. We also highlight that if the observables do not fully span the support of the Green's function, say due to a truncated dimensionality, we will lose information, and this could produce artifacts in the final reconstruction.

7.3 Experiment

7.3.1 State Preparation

In practice, to generate our temporal modes, we carve pulses from a CW diode laser (Sacher, Littrow ECDL), which is passed through an integrated fibre Mach–Zehnder interferometer (FMZI), which enables arbitrary control of the transmitted amplitude and phase, via electro–optic modulation of the phase in each arm. Each arm of the interferometer is driven by an arbitrary waveform generator (Tektronix AWG70000), amplified with a 10.7 Gb/s Driver Amplifier (PicoSecond Pulse Labs 5868) and the two are synchronised (Tektronix AWGSynch).

The AWGs were controlled by custom control software written in Python, which is discussed in appendix section A.3.1. Each of the amplifiers have independent control of the gain and crossing point offset, to account for discrepancies observed between the response of the two electrodes. The choice of settings for these are discussed in section A.3.1(a).

The optical output field was monitored during the course of the measurement using a 90:10 fibre beam splitter to pick off light to be sent to a fast photodiode (Thorlabs PDA-8GS) which was sampled by an oscilloscope (Lecroy Wavesurfer 10). To synchronise with the 80 MHz clock of the master laser (Tsunami Ti:Sa oscillator) which generated the control pulses for memory experiments, the AWG sample rate was restricted to 6.4 Gs/s (out of 50 Gs/s available). The frequency of the signal diode laser was logged with a wavemeter (High-Finesse/Ångstrom) throughout the course of the measurement.

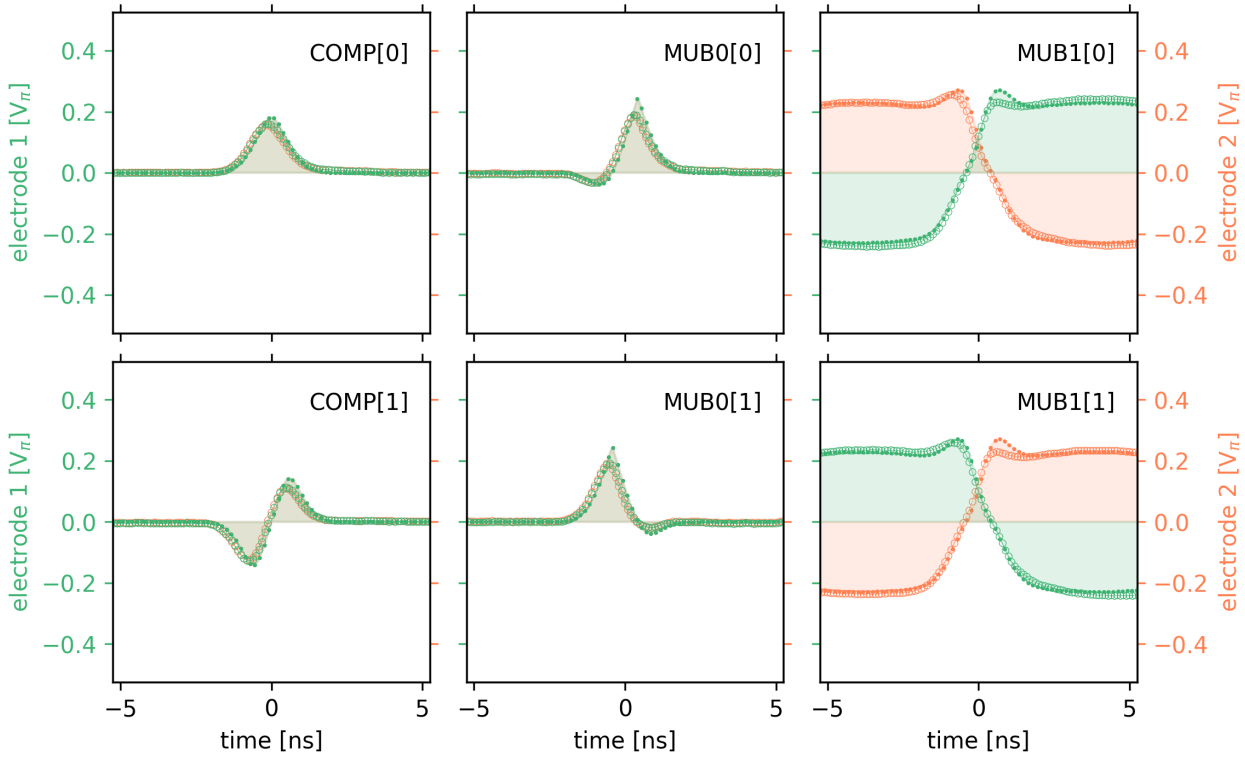
The carved signal incident on the memory is attenuated to contain $\langle N_{\text{in}} \rangle \sim 9$ photons. For the signal and control to be of commensurate bandwidth, the control field was passed through a 150 MHz etalon (5 GHz FSR) yielding a FWHM pulse duration (in intensity) of $\sim$ 600 ps. The limited sampling rate and amplifier bandwidth prevented the preparation of higher–order Hermite–Gaussian modes compatible with the control field bandwidth, and therefore restricted the present investigation to a two–dimensional temporal mode basis in which the signal states are written with FWHM of 1 ns.

In this case we wish to carve out a set of MUBs derived from a "computational basis" defined by the HG polynomials. These are constructed using the decomposition in [204].

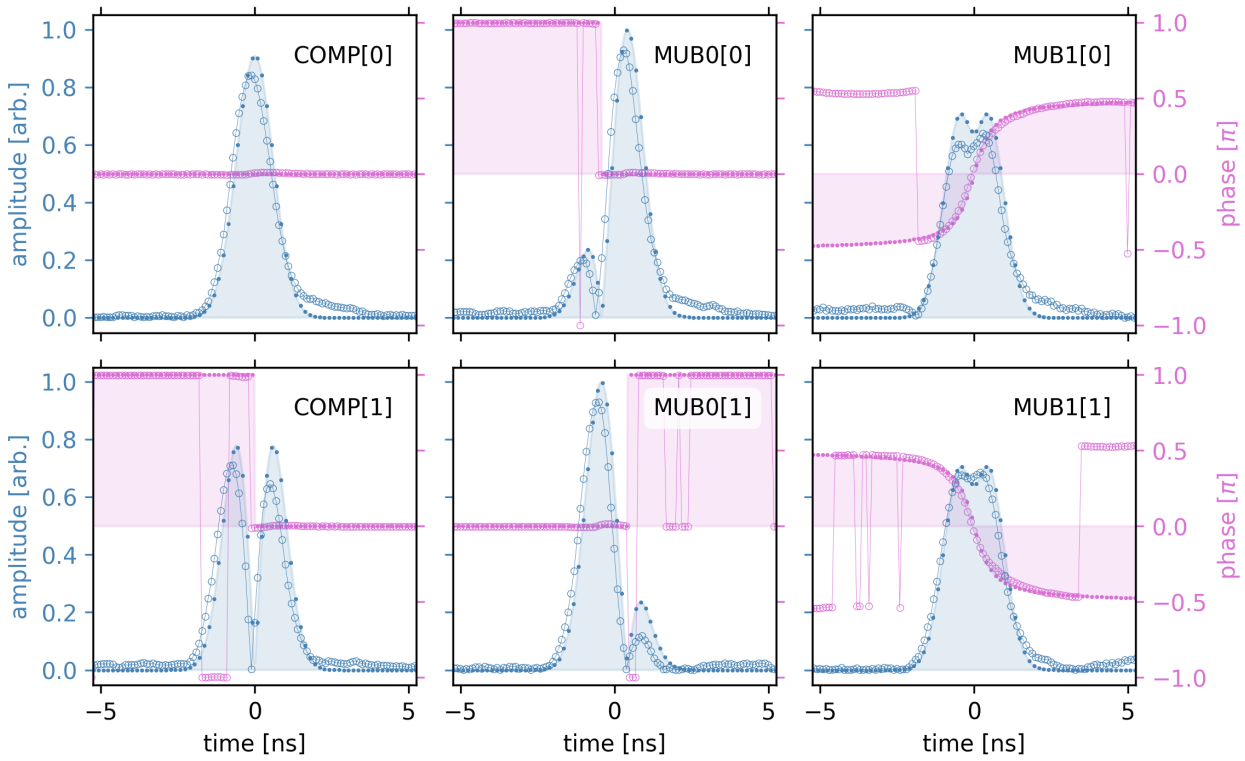
Starting from a desired signal field, with an amplitude and phase profile

$$E_{\text{s}}(t) = a(t)\mathrm{e}^{\mathrm{i}\phi(t)}, \tag{7.6}$$

we must configure the amplifier parameters such that the voltages written by the AWG map appro-



(a) Target (dots, shaded area) and measured (circles) amplifier output signals



(b) Target (dots, shaded area) and *inferred* (circles) output optical field

Figure 7.1: Preparation of MUB states: (a) electronic amplifier output signals, (b) resulting optical field from inverting equation 7.7. The shaded area shows the target profile, dots show the values written to the AWG (i.e. the target signal sampled at the AWG sample rate), and the circles signals derived from direct measurement of the amplifier output.

priately[5] into the output optical field:

$$V_1(t)/\alpha = \arcsin[a(t)] + \phi(t),$$
$$V_2(t)/\alpha = \arcsin[a(t)] - \phi(t), \tag{7.7}$$

where V_i for $i \in \{1, 2\}$ are the amplifier output voltages, and α is some constant of the EOM. Denoting the AWG outputs voltages as U_i applied to each amplifier, the mapping $U_i \to V_i$ depends on the particular settings on the amplifier. To minimise distortion of the output pulse, we must ensure that we operate within the amplifier's linear response regime, such that $V_i = \gamma U_i$. Note also that (a) the frequency response of the amplifier; and (b) the sign dependent response; must be considered.

In order to verify the optical signal produced, there exists a well established toolbox of reconstruction methods in ultrafast optics (see [205] for a review of pulse characterisation methods) for ∼THz bandwidth pulses, since these are based upon spectral measurements. Interactions in atomic physics operate in a different regime. To enable the characterisation of more narrow–band light, with bandwidths in the MHz–GHz regime and hence compatible with atomic memories, we have developed and demonstrated a novel pulse characterisation technique *Temporal Electric-field Amplitude + Phase: Algorithmic Reconstruction via Time-domain Interferometry* (TEA+PARTI), which is presented in appendix B.

However, due to the components available for our experimental demonstration of TEA+PARTI, the infrastructure is restricted to a maximum pulse bandwidth of $\lesssim$ 400 MHz (see appendix B), whereas in this experiment we are constrained to signal pulses $\gtrsim$ 1 GHz to remain compatible with a control field derived from the Ti:Sa. However, we anticipate that TEA+PARTI will be a valuable tool for the future development of this experiment.

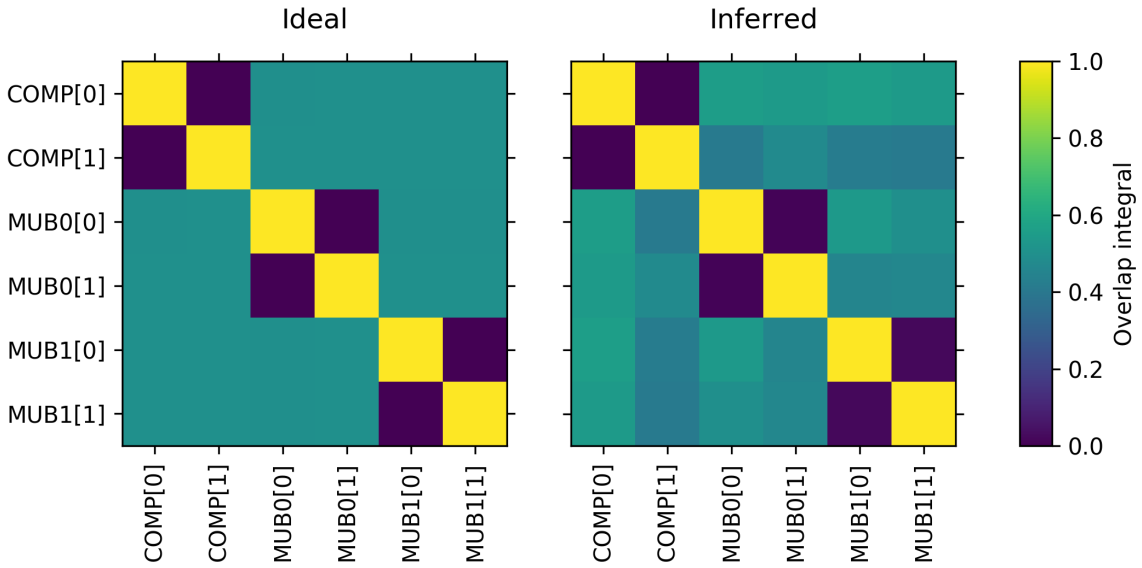
Without direct optical characterisation possible at present, therefore, we estimate the optical signals inferred from measurement of the voltage signals applied to the carving interferometer. After calibration of the gain,[6] the voltage signals and the correspondingly inferred optical fields, measured for the three mutually unbiased bases (in dimension two) for our experiment are shown in figure 7.1. In this case we still observe some imbalance between the electrodes and some washing out of the finer features.

This calibration is not perfect which motivates the development of tools for the optimisation of electrode voltages *for each signal*, which is investigated in appendix C, which, again may be a useful tool for future implementations, particularly when accessing a higher dimensional temporal–mode basis, where dynamic tuning of the electrode signals can help compensate for ringing. Ultimately it would be appealing to optimise based upon the optical signal – that is – defining a cost function derived from the optical characterisation of the signal, rather than the target electrode voltages.

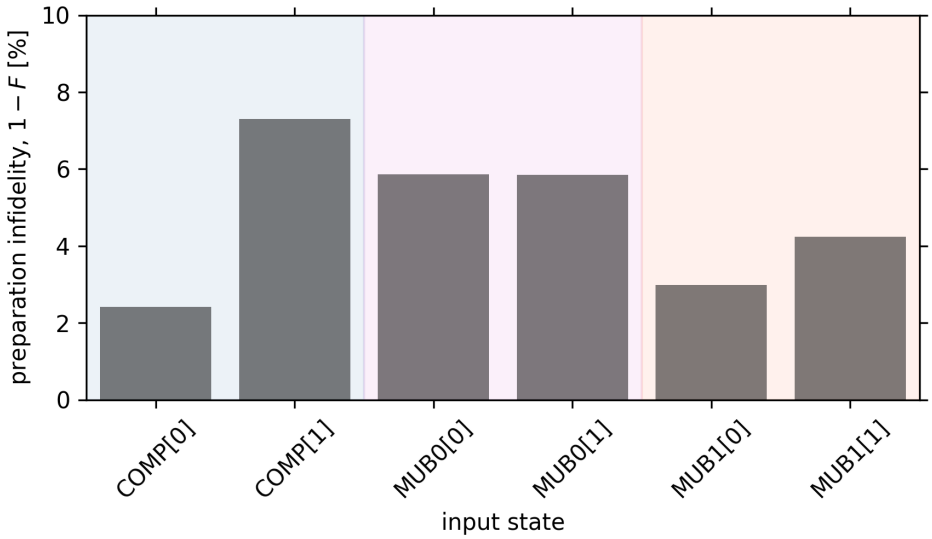
With the optical fields inferred from the electrode voltages, an *optimistic* estimate of our state preparation fidelity, and the mutual overlaps are summarised if figure 7.2, where we estimate the fidelity between two pulses as $F_{ij} = |\int \mathrm{d}t' E_i^*(t') E_j(t')|^2$ where the indices i, j label both the state identity and denote target or inferred field.

Upon measuring the intensity of the output fields, we observe the traces shown in figure 7.3. Whilst this mismatch may in part be explained by including a convolution with the detector response causing the optical signals to be "smeared out" in time, the fields inferred from the electrodes do

[5] See appendix B for a more detailed derivation.
[6] Such that the peak voltage optical signal is written to span the V_π of the electrodes.



(a) Mutual overlap of prepared MUBs



(b) Metric of fidelity with target state

Figure 7.2: (a) Mutual overlap of inferred probe state fields. (b) Overlap of inferred field with target states. Whilst we do not have direct access to the optical fields in these cases, TEA+PARTI yields a ∼70% agreement between the inferred fields and the direct reconstruction of the optical fields.

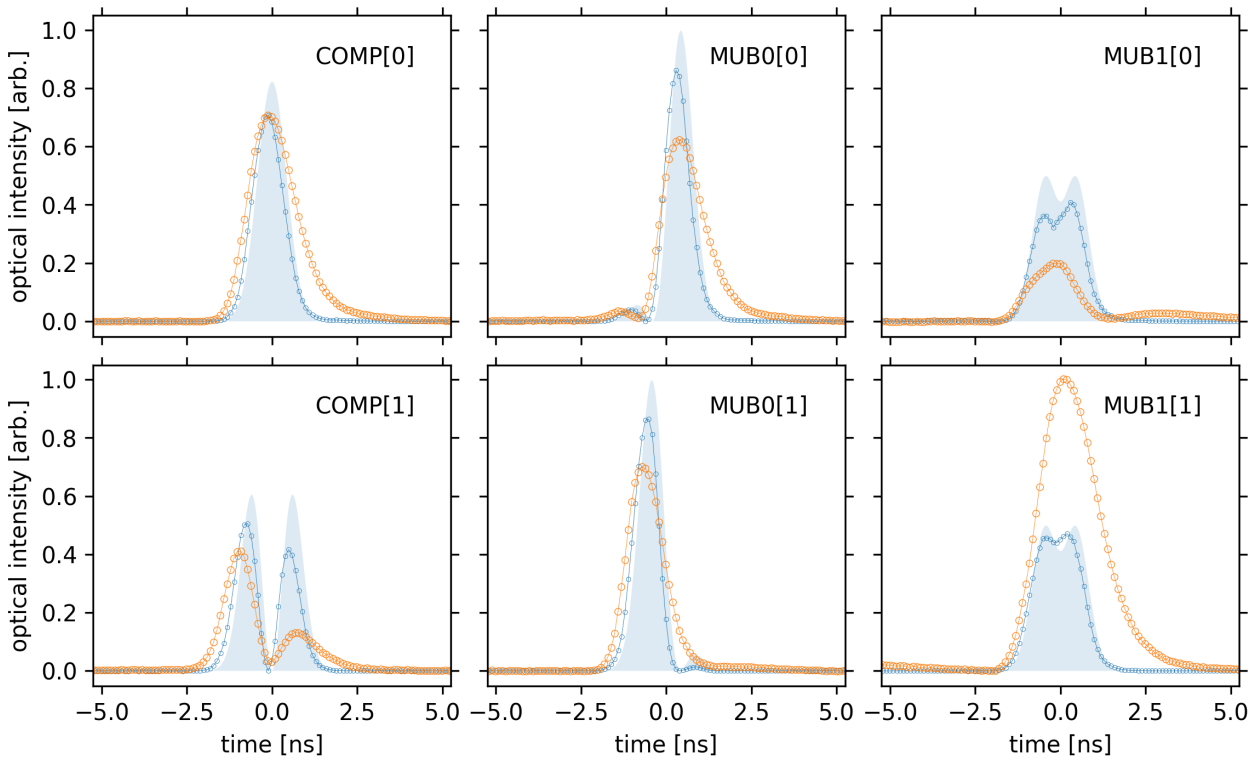


Figure 7.3: Optical output intensity: shaded area indicates the target profile, blue as inferred from electrode inputs and orange directly measured from photodiode signal. .

not capture the full picture. Based upon TEA+PARTI for longer pulses in appendix B, we observe some disagreement between the directly reconstructed optical state and that which is inferred from measurement only of the electrode voltages, typically exhibiting an overlap $\gtrsim 0.7$ between the two.

One contributing factor limiting the ability to infer the output optical field is the bias voltage of the FMZI which was observed to drift over the timescale of an experimental run. To constrain this during memory measurements, with a repeatable procedure, we periodically minimise the CW leakage prior to the carved signal using the observed single photon counts.

A further observation from the optical output fields in figure 7.3, in comparison with the written electrode fields in figure 7.1, is that the COMP and MUB0 bases are both symmetric in the written electrode fields, whilst the two MUB1 states are of opposite sign. This sign difference may account for the asymmetric distortion seen in the optical output in MUB1 $|0\rangle$ and $|1\rangle$ fields, as this would be particularly sensitive to the bias.[7]

7.3.2 Memory infrastructure: Modifications

The general scheme of the Raman memory experiments conducted here was similar to that presented in chapter 3. In this section I outline the improvements made for this experiment, which are based upon some of the results in chapter 4, along with some infrastructure upgrades. A sketch of the modified schematic is shown in figure 7.4.

For the purposes of these demonstration measurements, we operated with a free-space Raman memory using a vapour cell length of 7.5 cm, containing 5 Torr N_2 buffer gas. In the absence of the cavity, the optical pumping efficiency w was more straightforwardly determined by scanning the CW

[7] Although the bias, in isolation, is not sufficient to capture this asymmetry – since the transmission through the FMZI at times well separated from the signal pulse is minimised.

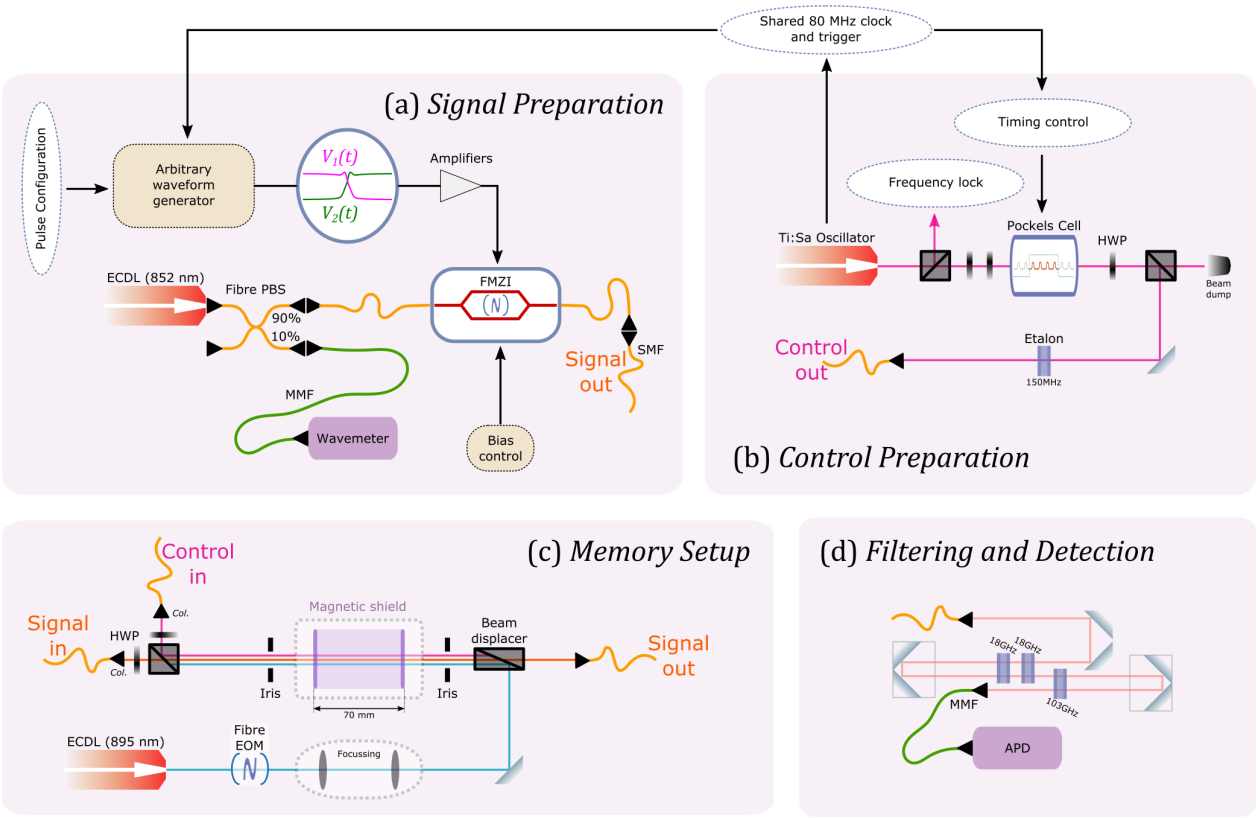


Figure 7.4: Schematic for the modified experimental elements for interfacing carved signal pulses with the Raman memory.

diode laser across both of the atomic resonances.[8] Likewise, the optical pumping was performed using a seperate diode laser (Toptica DL100) now locked to resonance with the D_1 line. As discussed in chapter 4 the move to a molecular buffer gas, in conjunction with pumping on the D_1 line, enables high pumping efficiencies to be reached [142].

Although for the present analysis we will only consider the *storage* efficiency of different signals, retrieval was also performed. Some typical memory traces for these measurements are shown in appendix E. Due to the optical noise observed following the carved pulses,[9] memory measurements were performed with pulses picked with a delay of 112.5 ns between storage and retrieval pulses. Because of the longer storage times the pump light was passed through an intensity modulator (EOSpace) in order to *quickly* switch off the pump light for the storage duration, with an extinction 100:1. The bulk cell was enclosed in μ–metal, in the same configuration as [142] suppressing the residual field by a factor $\sim 10^3$, to avoid magnetic dephasing impacting the memory lifetime.

The Pockels cell was also upgraded (Leysop, RTP) with a new driver (BME Bergmann) now allowing a repetion rate of 2 MHz, compared to 6 kHz in the previous cavity demonstration. The increased repetition rate is enabled by a reduced piezoelectric constant, this improvement reduces the time necessary for the acquisition of counts on the APDs, but also means the response is less susceptible to ringing and results in an improved extinction ratio of 20,000:1 which minimises leakage of the control into the memory preceeding the memory interaction window.

The detection stage for these measurements was also modified from the previous cavity imple-

[8] For a typical absorption spectrum, see figure A.4.

[9] All fibres between the signal preparation and the memory were APC to prevent back-reflection/"etaloning" – this noise is therefore most likely leaked diode light due to ringing of the amplifiers.

mentation.[10] In this case we operate with two double passed 18.4 GHz etalons, followed by a double passed 103 GHz etalon, before being coupled to a single APD (Perkin Elmer). The transmission from the memory to the detector was $\sim 3\%$.

In these measurements we also switched to a detuning of -15.2 GHz, and pumped into the F=3 ground state level. In this configuration the (previously defined) A field involved in FWM is now lower in energy than the signal (and the terminology "anti–Stokes" is no longer strictly applicable). This configuration was adopted due to the Ti:Sa becoming unstable at the usual frequency.

In this chapter, we seek to characterise the $\Lambda-\Omega$ system, and thus probe it with a variety of different states, and record the measured "read–in" efficiency. We record the memory efficiency for the six MUB states in two dimensions, and for reasons that will be motivated in section 7.4.3 we perform measurements for each of these at a series of time delays of the signal arrival time, with respect to the control pulse. The delay is applied electronically by timeshifting the waveform written by the AWGs with respect to the trigger recieved from the Ti:Sa.

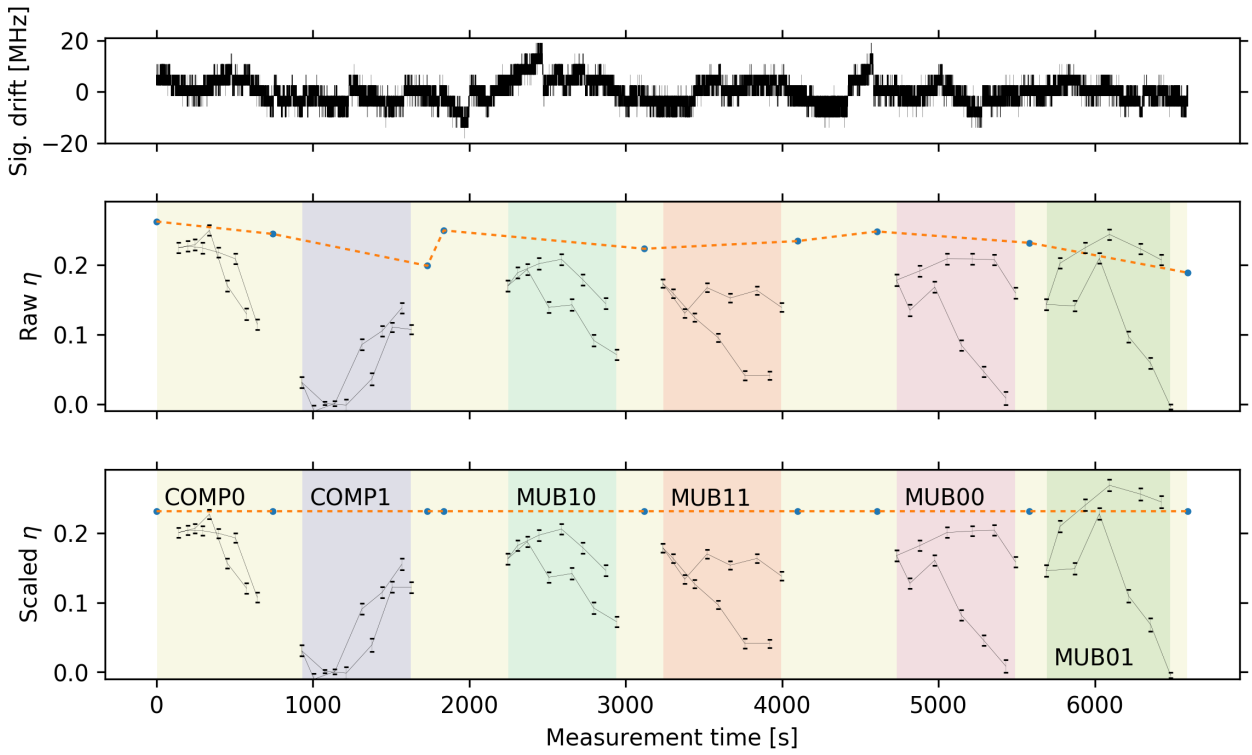


Figure 7.5: Top: frequency of the diode laser; middle: the raw storage efficiencies extracted; bottom: rescaled efficiencies. For each probe state, the points are joined in order of their time delay with respect to the control pulse – we begin each set at zero delay, and oscillate the delays ($+\,\mathrm{d}t$, $-\,\mathrm{d}t$, $+2\,\mathrm{d}t$, ...). The drift in efficiency for the reference state (blue circles) do not directly follow the diode drift, and are also subject to other experimental factors which are not directly accessible, such as Ti:Sa, optical alignment, detection filtering.

Finally, we note that over the course of a measurement, many of the experimental components suffer drift, for instance the laser frequency and the FMZI bias voltage. We interleave measurements with a standard reference measurement of the COMP0 state at zero delay with respect to the control field, so that in order to compare the memory efficiencies in each case, we then estimate the scaled efficiencies for a given state at a particular delay, by assuming a *linear* temporal drift between reference measurements. This is depicted in figure 7.5, where the top trace shows the signal frequency offset

[10] These are under constant development in the lab, as they are shared with multiple experiments.

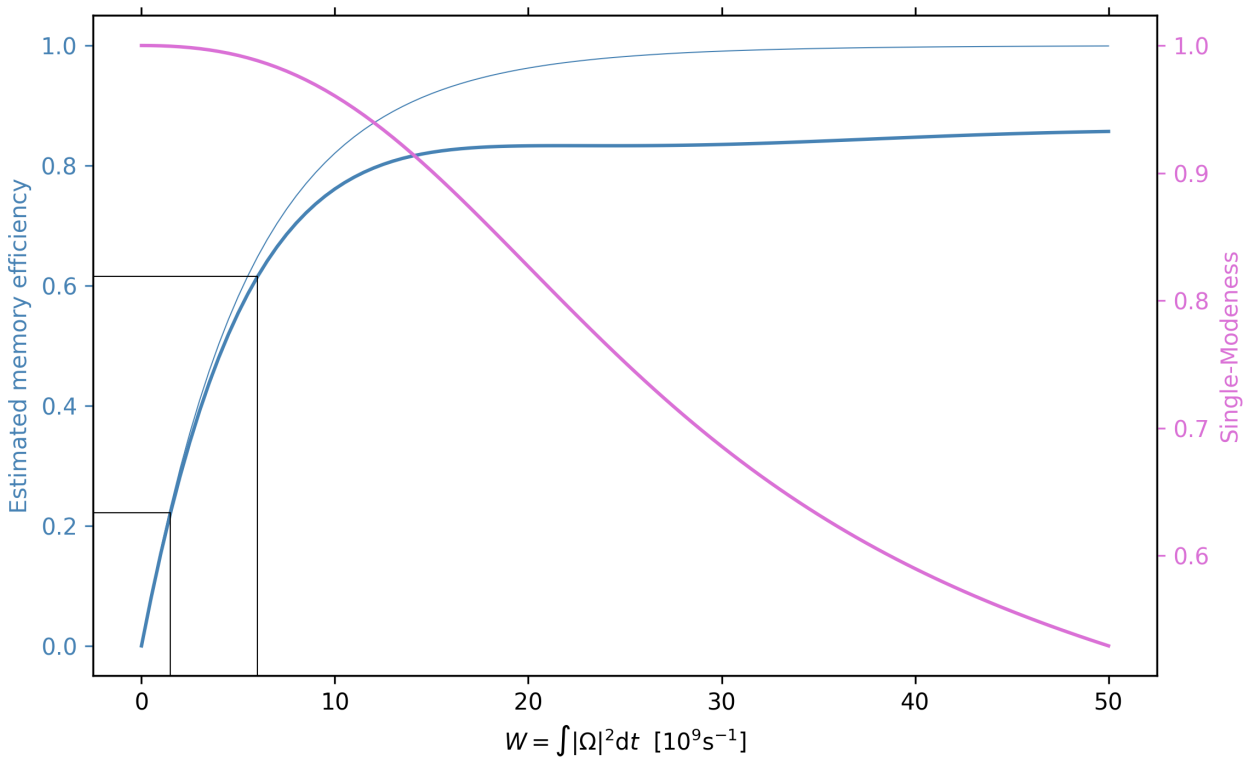


Figure 7.6: Model of memory performance, for a Gaussian control pulse, with varying pulse energy. The heavy blue line shows the estimated efficiency given a signal in the inferred "COMP0" mode, whilst the thin blue curve shows the efficiency if the signal is sent in the principle eigenmode of the memory, and (in this model) can reach unity. The purple curve on the right axis shows a metric of single-modedness, described in the text. The horizontal lines correspond to the measured COMP0 efficiency.

from the target frequency over time as monitored on the wavemeter, the central trace plots the raw recorded memory efficiency for each state/delay and the lower trace shows the same efficiencies after rescaling all the efficiencies assuming a linear drift in time between the reference measurements.

7.3.3 Preliminary Conclusions

Before proceeding to a more in depth tomographic analysis of the obtained efficiencies, let us first review the direct output of the memory measurements.

Figure 7.6 shows the predicted memory efficiency of the COMP0 reference state using the simplified model in equations 7.1. We show the predicted efficiencies (heavy blue line) assuming an input with the inferred COMP0 profile sent to the memory synchronously with the control, and optimal efficiency (light blue line) if an input in the principle eigenmode of the memory is prepared, as the amplitude of the Rabi–frequency is scanned. These are plotted for a fixed $d\gamma = 1.6 \times 10^{12}$ s^{-1} and Gaussian control pulse with a FWHM of 590 ps, where both of these are estimated from the experiment.[11] The two horizontal lines mark the measured efficiency for reference COMP0 measurements at two control field powers. On the right hand scale, I show the single modedness of the Green's function as defined in equation 7.5.

Figure 7.7 shows the measured memory efficiency as a function of delay between signal and control for the 6 MUB states constructed in two dimensions, at two different control field powers. The solid

[11] This can be used to obtain a secondary estimate of W.

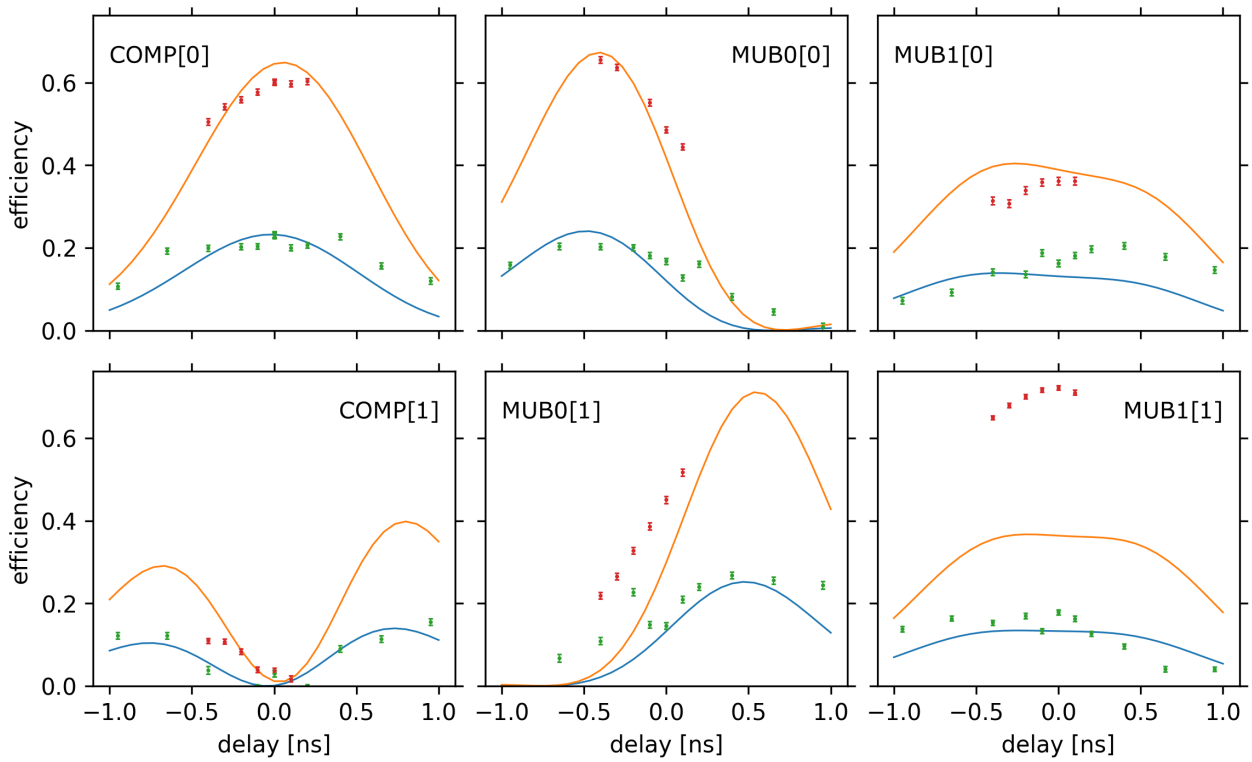


Figure 7.7: Measured read-in efficiencies for the timescanned MUBs in dimension 2, with *no free parameters* – these correspond reasonably well with estimated experimental parameters. The two series are conducted with different control field powers: blue/green at low power and orange/red at $\sim 4\times$ control pulse energy. Error bars on the data derived from Poissonian noise.

lines are models for the predicted efficiencies, given our prepared signals – inferred from the electrode inputs – with our estimates of $W = 2 \times 10^9\ \mathrm{s}^{-1} \times \{1.0, 4.2\}$ for the two data sets. The two lines are constrained to have the same pulse duration and optical depth, and an appropriate factor between W corresponding to the ratio in control field power measured for each case.

We know that the model is incomplete and doesn't fully capture the efficiencies seen in figure 7.7. At these control field powers it is unlikely that this is due to FWM, and we notice the corresponding asymmetry in height between the MUB1 states follows that observed in the output optical fields in figure 7.3.

Nevertheless, it is informative to employ this toy model for the purposes of benchmarking tomography on a known density matrix and efficiencies that derive from it.

From the functional form of these data, in which we perform time scans of the signals with respect to the control, the storage efficiency approximately maps out the input intensity of the signal convolved with the control pulse. We notice, for instance, from the read-in efficiency of the COMP1 state (which we expect to be approximately orthogonal to the control pulse) that the storage efficiency at zero delay is suppressed. Functionally we also note that as the power is increased there is some distortion, particularly noticeable in the COMP1 efficiency, where the central zero is shifted in time and no longer reaches zero efficiency.

7.4 Reconstruction of the Green's Function

Based upon these measurements of the memory efficiency, we seek to use this information to gain insight into the structure of the memory interaction. This will enable us to investigate the applicability of the Raman memory as a mode selective device.

To extract this information we aim to tomographically reconstruct the matrix which describes the memory interaction, and from the reconstructed state obtain information about the modal structure of the process.

Firstly, in section 7.4.1, I outline the maximum-likelihood method of quantum state estimation and the techniques used. Then in section 7.4.2, I discuss the general way in which this can be applied to the memory process – and from this extract how the A−Ω system interacts with temporal modes; and then in section 7.4.3, consider this mapping in practice, and discuss how it relates with our experiment.

7.4.1 Informationally Incomplete Tomography

In order for us to reconstruct the Green's function of the memory interaction we first outline the necessary methods from quantum state tomography employed here. Given *many copies* of some unknown state with a density matrix ρ, upon each of which we may perform a set of projective measurements $\{\Pi_j\}$; we wish to reconstruct ρ based upon the observed measurement results. If we consider a set of such measurements, the probability of observing an outcome j is

$$p_j(\rho) = \mathrm{Tr}\left[\rho\Pi_j\right]. \tag{7.8}$$

Extended to an ensemble of N measurements, the *likelihood*, $\mathcal{L}$, of observing some set of results, given the underlying density matrix ρ can be expressed by the product

$$\mathcal{L}(\rho) = \frac{N!}{\prod_k (n_k!)} \prod_j p_j(\rho)^{n_j}, \tag{7.9}$$

where the prefactor is an unimportant normalisation constant, and $f_j = n_j/N$ are the relative frequencies (i.e. $\sum_j f_j = 1$).

In order to find the ρ that maximises the likelihood function, we implement an iterative algorithm starting from an unbiased estimate of ρ, and propagate this by diluting it with an operator constructed from the projectors and observed frequencies which evolves towards a state that maximises $\mathcal{L}$.

We adopt the diluted maximum-likelihood method given in [206, 207]. Details of the proceedure are provided in appendix D, and we quote the main results here without proof. The essence of the method is to construct the operator

$$\mathbb{R} = \sum_i \frac{f_i \Pi_i}{\mathrm{Tr}\left[\rho\Pi_i\right]} = \mathbb{1}_\rho, \tag{7.10}$$

for the unknown, extremal state ρ, and set of measurements $\{\Pi_i\}$ and observed outcome frequencies f_i. where this may be interpereted as a closure relation, where $\mathbb{1}_\rho$ represents the identity operator on the support of the extremal density matrix which may be reconstructed *given the set of measurements* $\{\Pi_i\}$.

Note that in general this closure relation in $\mathbb{R}$ may only span a subspace of the full space of the

true density matrix, where a set of informationally incomplete measurements do not have access to components orthogonal to this space [208, 209].

Proceeding with the algorithm we take a small positive dilution parameter ε, and define a modified propagator:

$$\mathbb{R}' = \frac{\mathbb{1} + \frac{\varepsilon}{2}\mathbb{R}}{1 + \frac{\varepsilon}{2}},$$

and so at some iteration k of the dilution, the next estimate for ρ is given by

$$\begin{aligned}\rho^{(k+1)} &= \mathbb{R}'\rho^{(k)}\mathbb{R}' \\ &\approx (1-\varepsilon)\left[\rho^{(k)} + \frac{\varepsilon}{2}\left(\mathbb{R}\rho^{(k)} + \rho^{(k)}\mathbb{R}\right) + \mathcal{O}\left(\varepsilon^2\right)\right] \\ &= (1-\varepsilon)\rho^{(k)} + \varepsilon\bar{\rho}^{(k)},\end{aligned}$$

i.e. the next constructed from a *convex mixture* of $\rho^{(k)}$ and $\bar{\rho}^{(k)}$, where we have defined

$$\bar{\rho}^{(k)} = \frac{1}{2}\left[\mathbb{R}\rho^{(k)} + \rho^{(k)}\mathbb{R}\right]. \tag{7.11}$$

This approach then provides us with a straightforward recipe for the maximum likelihood reconstruction of some density matrix, based on an ensemble of imperfect measurements available.

7.4.2 Mapping to Memory Experiment (1): Theory

To relate this procedure to the problem of reconstructing the Green's function of the memory interaction: if we consider a single trial of the memory, having a single photon incident on the system. The read–in efficiency quantifies the probability that the photon is written into the memory. We may relate this to the POVM-like picture, whereby in each trial, we have two outcomes: *success* ("photon stored") or *failure* ("photon transmitted"). In each trial, we choose some measurement settings: prepare a probe signal in a particular temporal–mode and arriving at the memory with some temporal delay with respect to the control pulse. Over N trials, we accumulate statistics for the memory efficiency, from which we extract the "frequency of success".

We seek to characterise the Green's function which maps an input temporal mode into the memory:

$$\begin{array}{ccc} B(z) & = \int \mathrm{d}t \left[\quad K(z,t) \quad S(t) \quad \right] \\ \begin{pmatrix} \{z_j\} \\ \downarrow \end{pmatrix} & = & \begin{pmatrix} & \{t_i\} \rightarrow \\ \{z_j\} \downarrow & \end{pmatrix} \begin{pmatrix} \{t_i\} \\ \downarrow \end{pmatrix} \\ \boldsymbol{B} & = & \boldsymbol{K} \quad \boldsymbol{S} \end{array} \tag{7.12}$$

The matrix $K(z,t)$ encapsulates the A$-\Omega$ system, via equation 7.1. This may be equivalently expressed in the form $\boldsymbol{K} = \sum_i \lambda_i |\beta_i\rangle \langle\sigma_i|$, in that it acts on an optical mode $|\sigma_i\rangle$ and maps this into a spin–wave mode $|\beta_i\rangle$. The singular values λ_i appearing here constitute the weights of the respective eigenmodes of the memory.

If we then probe this with a state $S_\alpha(t)$, in which the subscript α encodes the mode and time delay of the probe, the probability of storage (i.e. the number of spin–wave excitations created) is

$\int \mathrm{d}z B^\dagger(z)B(z)$ or $\boldsymbol{S}^\dagger\boldsymbol{K}^\dagger\boldsymbol{K}\boldsymbol{S}$. If we therefore define:

$$\begin{aligned}\boldsymbol{\rho} &= \boldsymbol{K}^\dagger\boldsymbol{K} \\ \boldsymbol{\Pi}_\alpha &= \boldsymbol{S}_\alpha\boldsymbol{S}_\alpha^\dagger/\mathrm{Tr}\left[\boldsymbol{S}_\alpha\boldsymbol{S}_\alpha^\dagger\right],\end{aligned}$$

then,[12] the read–in efficiency is simply $\eta_\alpha = \mathrm{Tr}\left[\boldsymbol{\rho}\boldsymbol{\Pi}_\alpha\right]$, analogous to the likelihood expressions previously. In the above decomposition, we have $\boldsymbol{\rho} = \sum_i |\lambda_i|^2 |\sigma_i\rangle \langle\sigma_i|$, and $|\lambda_i|^2$ is (related to) the efficiency with which the eigenmode $|\sigma_i\rangle$ is written into the memory.

The definition ρ is consistent with the properties of a density matrix: that is

1. Hermiticity: By construction this is satisfied $\boldsymbol{\rho}^\dagger = \left(\boldsymbol{K}^\dagger\boldsymbol{K}\right)^\dagger = \boldsymbol{K}^\dagger\boldsymbol{K} = \boldsymbol{\rho}$;
2. $\mathrm{Tr}\left[\boldsymbol{\rho}\right] = \sum \lambda_i$, with $\eta_i \propto \lambda_i^2$. Without loss of generality, we are free to set the "gain" such that this is 1.

In principle, the dimension of the TM space is of course infinite; however due to experimental limitations, we may only have access to part of this space.

7.4.3 Mapping to Memory Experiment (2): Practice

We now develop these ideas and consider how they may be applied, taking into account the specifics of the present experiment. Owing to the control–pulse bandwidth and sampling rate of the AWGs, in the present experiment we are restricted to a two dimensional subspace. Thus our *informationally incomplete* measurement may yield a trace of less than unity within this domain. The trace, however *is constrained*[13] and as above, we may set this to one.

Further to this, we also recall our imperfect state preparation: in the first instance our signals are more narrow band than the control field, by a factor ~ 2, but also suffer imbalance and ringing of the electrodes. Nevertheless, provided that the states are known, and have non–zero overlap with the ideal MUBs, we should have some access to the phase information needed to fully reconstruct the memory interaction.

To attempt to compensate for these we expand the probe space by including *time–scans* of the input states by electronically delaying the generation of the signal. We anticipate that, in contrast to having single time measurements, these linearly independent states in principle enable us to access more information about the temporal profile of the interaction. Although this decomposition will not be optimal as compared with a higher dimensional ensemble of MUB states.

Alternatively, in future developments of this experiment, with improvements to the infrastructure we hope to be able to gain access to a higher dimensional space, in which case this additional time information may not be necessary, since with sufficiently high dimension, the HG bases provide a complete basis into which an arbitrary function may be decomposed.

If we can prepare an ensemble of these probe states, with a range of delays, A_α, we can probe the structure of ρ. Performing an eigenvalue decomposition on the reconstructed density-like-matrix, we seek to gain insights into the mode structure of the memory interaction.

For both experiment and simulation, we may take the inferred fields (with appropriate delay) to construct the projectors, and the corresponding efficiencies as the frequencies to pass to the tomo-

[12] Including the denominator in this expression $\mathrm{Tr}\left[\boldsymbol{S}_\alpha\boldsymbol{S}_\alpha^\dagger\right]$ is equivalent to normalising by the input state excitation number $\int \mathrm{d}t S^\dagger(t)S(t)$.

[13] And the Lagrange multiplier method in appendix D holds, regardless of the actual value between 0 and 1.

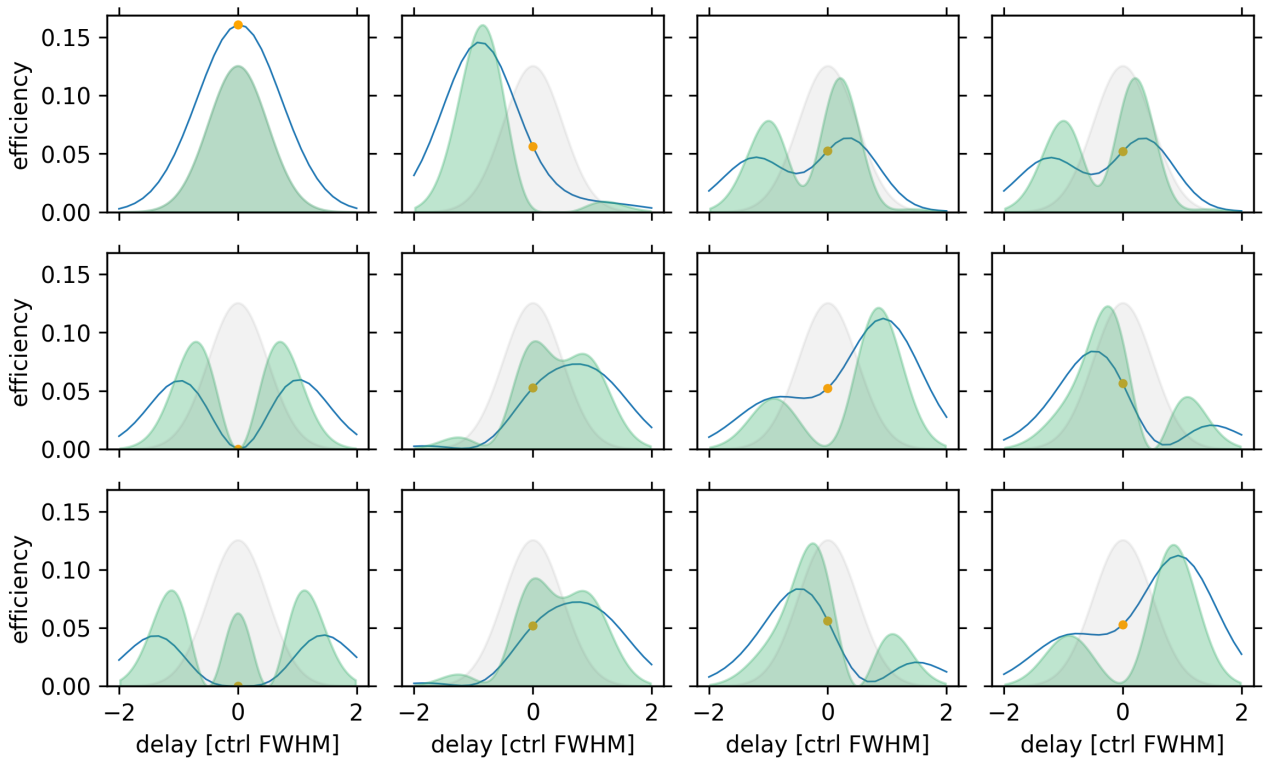


Figure 7.8: Simulated efficiencies for bandwidth matched signal and control for timescans of the mutually unbiased bases in three dimensions, in *low coupling* regime. Simulated efficiencies displayed on the blue curves, the shaded areas denote the *zero delay* optical intensities for the signal (green) and control (grey) pulses. The orange dots highlight the projected efficiencies at zero delay.

graphy algorithm. We pass the data shown in figure 7.7, along with the simulated efficiencies corresponding to the same experimental parameters.

In this approach, our projectors are constructed from the time trace of the signal electric field, whereby we interpolate the value at the abcissas and scale by the quadrature weights. The matrix ρ approximates a continuous density matrix which is represented directly in the time domain. To reduce the dimensionality of the problem, we then perform another coordinate transformation mapping the time domain onto an orthonormal basis set, further discussion of this choice is given in appendix E.

7.4.4 Performance Benchmarking

To investigate the performance *and limitations* of the application of this algorithm to our experiment, we may consider the estimated memory efficiencies that one would obtain given a fully known interaction map – assuming an *ideal* experiment, i.e. perfect knowledge of the input states, and measurement results. For this we use our simplest model of the memory interaction given by 7.1 and simulate the efficiencies for some K given a set of *known* S_α. Plugging these states and simulated efficiencies should allow us to gain insights into the algorithm performance.

Furthermore, given our restricted space, we can use simulation to investigate the consequences of experimental imperfection, such as the mismatch between signal and control, or the truncation of the probe state basis. For instance we also investigate being able to access higher dimensional space with more optimal measurements – say, with no time scan information, but higher dimensional MUBs – the (ideal) MUB states and resulting efficiencies for HG states up to $d = 3$, and the corresponding MUBs, are shown in figure 7.8.

Note that, in addition to experimental considerations, one must also be conscious of the choices of numerical parameters such as the dilution parameter ε for the algorithm, where in this case we balance numerical stability with computational overheads. For the results presented in the following section, we slowly vary ε from 10^{-3} to 10^{-5} as the resulting density matrix converges to some value.

We monitor the evolution of ρ as the algorithm progresses, we record the trace distance[14] between successive iterations of ρ, as well as the overall distance from the initial (unbiased) state. We also record a metric of the "likelihood" that the observed frequencies could be derived from ρ given the form of our measurements. By all of these metrics, the algorithm generally seems to converge, and successfully exits once the rate of change of these falls below some tolerance. See appendix E for further discussion on the choice of these parameters.

The tomography algorithm doesn't *a-priori* know about the magnitude of the efficiency, and just relies upon relative values, so without loss of generality we constrain that the trace is 1 in each step of the algorithm, since the magnitude of the efficiency just provides the "gain" of the process,[15] and determines the trace – so, of itself, the trace distance is a little bit arbitrary.

7.5 Results

We discuss a preliminary, broadly qualitative, analysis of the performance of this reconstruction method on our experimental data (§7.5.1). Then, we move on to explore potential alternative measurement configurations via simulation (§7.5.2), for measurements not accessible with the present experimental infrastructure. Some discussion of the insights and consequences of these is given in section 7.5.3 to guide the direction of future experiments.

As a first step, we look at the eigenmodes, and eigenvalue ratios compared between those obtained from (i) the reconstructed matrix we find through the tomography algorithm, and (ii) the simulated interaction matrix, directly.

Whilst investigating the mode structure of the interaction, we draw comparisons with studies of metrics of entanglement. In this case, it is common-place to consider the Schmidt decomposition in which the modes are distributed with weights λ_i, such that $\sum_i \lambda_i^2 = 1$. The effective mode number of the state is then given by $k = 1/\sum_i \lambda_i^4$ [152, 210].

We constrain the trace to be unity, so our eigenvalues of $\boldsymbol{K}^\dagger \boldsymbol{K}$ thus sum to one (and are related to the storage efficiency via $\lambda_i^2 \propto \eta_i$). It is expected that as the effective mode number of the interaction increases beyond the dimensionality of our probe space, that our reconstruction will break down as higher order basis states would be needed to access information and therefore reconstruct those modes.

7.5.1 Experimental Data

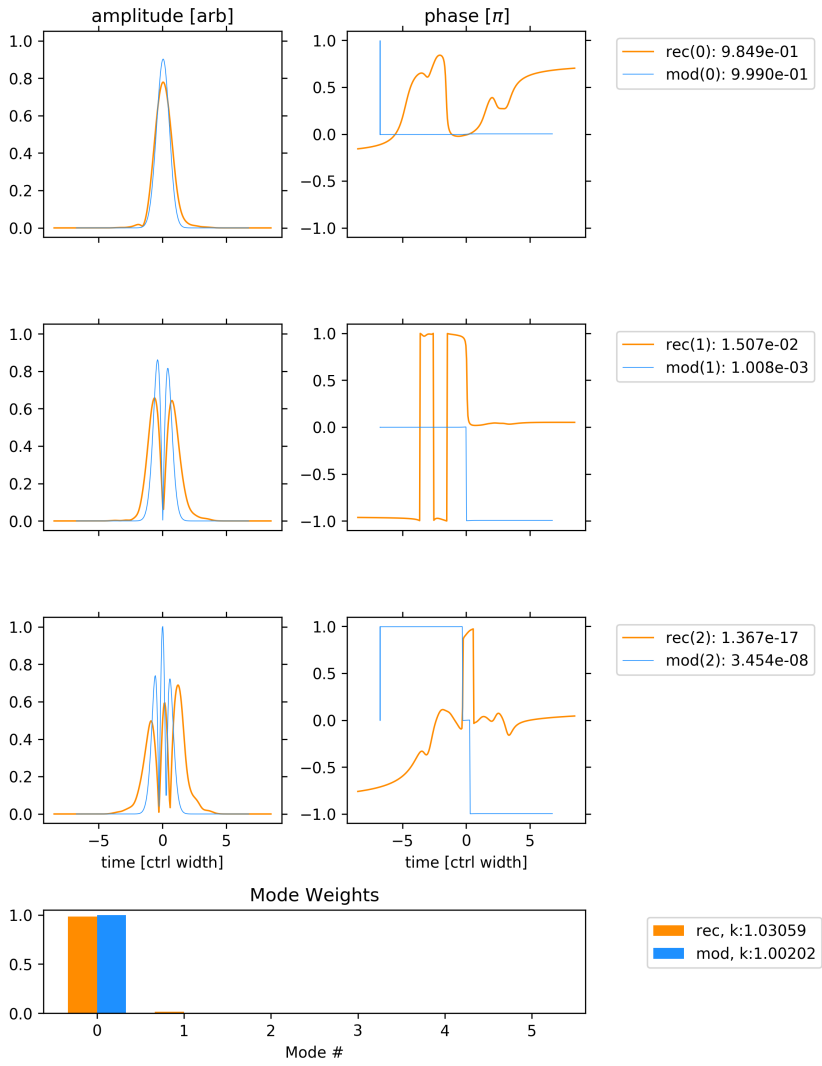
To graphically display the results of the reconstruction we consider the time-domain trace of the three principle eigenmodes of the simulated/reconstructed matrix, and their respective weightings, λ_i^2.

We first apply this method to the experimental memory data (orange) displayed in figure 7.7, and compare with the modelled memory eigenmodes (blue). These yield the modes plotted in figure 7.9.

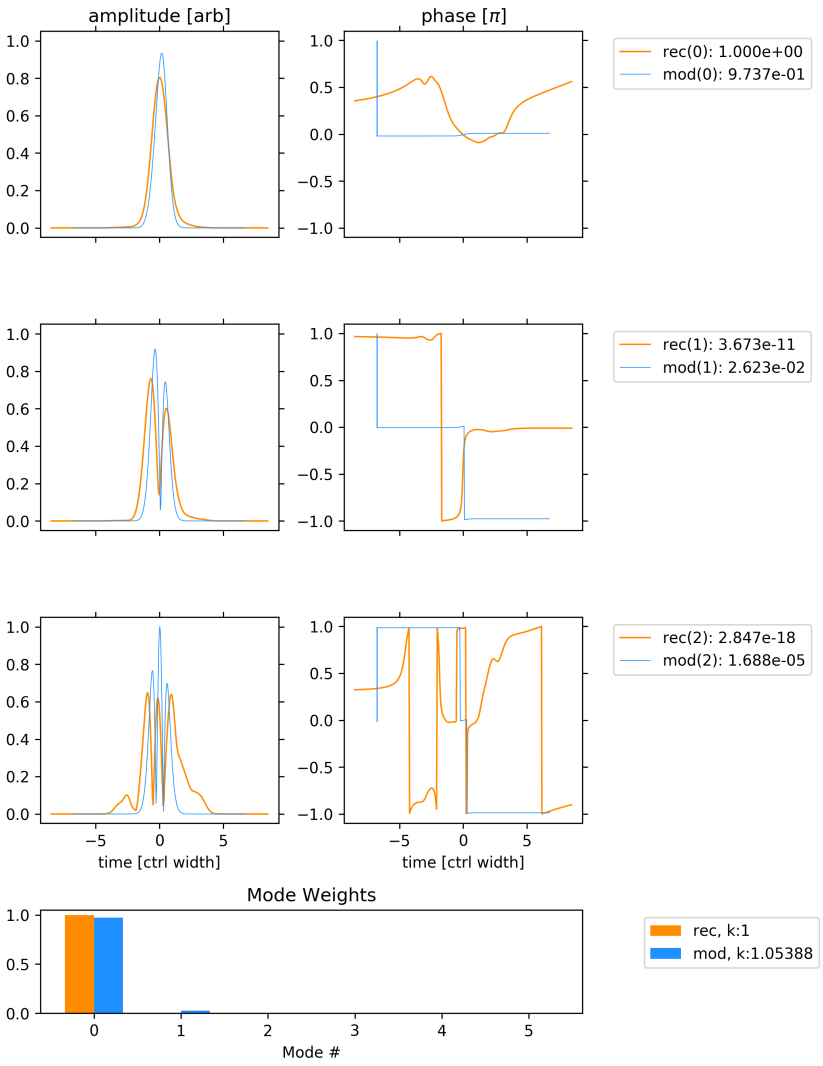
First, we look at the temporal shape of the reconstructed eigenmodes and note that these have inherited some features from the temporal profile of the input state – such as the tail on the falling

[14] If the trace distance is used to quantify the convergence of the algorithm, then the trace of the unknown matrix should be constrained to be constant.

[15] Since $\mathrm{Tr}\,[\boldsymbol{\rho}] = \sum_i \lambda_i \propto \sum_i \sqrt{\eta_i}$.



(a) Low power $W = W_0 = 2 \times 10^9\ \mathrm{s}^{-1}$



(b) High(er) power $W \approx 4\,W_0$

Figure 7.9: Reconstruction ("rec") on experimental memory data (orange). The blue traces "mod" display ρ derived from the models in figure 7.7. We show the time domain profile of the first three eigenmodes of the interaction matrix, along with their normalised eigenvalues. We observe good agreement at low gain – where we expect the simulation to be more reliable. In all cases the interaction is near single mode and the effective mode number $k \approx 1$.

edge of the pulse, and the resulting width is longer than predicted, given the control field profile (refer back to figure 7.1 showing the inferred probe fields). This is apparent for both powers.

These observations indicate that reconstruction via iterative dilution is sensitive to the profiles of the set of probe states, and determines the accessibility of width information.

We next turn to the mismatch in mode weightings, λ_i^2, and effective mode number, k. As noted the overall efficiency scaling in the algorithm is arbitrary – and since the higher power data has greater contrast between the $|0\rangle$ and $|1\rangle$ efficiencies,[16] the algorithm may tend to yield a *more single mode output.*

These discrepancies lead us to next look at the performance of the algorithm in an ideal case in which we fully know each part of the "experiment". That is, we have complete knowledge of (i) the underlying ρ, (ii) the exact shape of the probes Π_i, and (iii) the resulting efficiencies.

Nevertheless, from these data we already observe that certain features are well captured. For instance, the distortion/asymmetry of the second eigenmode as the control field power is increased. Focussing on the lower efficiency data, where we observe closer agreement between the model and experiment, we may confirm a high degree of mode selectivity of the interaction, and estimate a separability $\mathbb{s} \approx 98\%$ from the experimental data.

The uncertainty in this number comes principally from the restricted information available to the tomography. In the following sections we aim to investigate the limitations of the present experiment and how to maximise the information that may be extracted in future through our choice of measurement. If we may optimise on this, the next limitations to consider will arise from the ability to infer the form of the probe states which are passed to the algorithm.

7.5.2 Improving Measurement Choices

With these observations in mind, we next consider sending in probe states in a basis set which are matched in duration with the control field, and therefore in some sense more optimally span the support of the Green's function. Taking the *ideal* mutually unbiased input states in three dimensions synchronous with the control, we simulate[17] the efficiencies shown in figure 7.8 which are evaluated at comparable parameters for the experimental case. These yield the typical reconstructions shown in figure 7.10 for two different amplitudes for the control field.

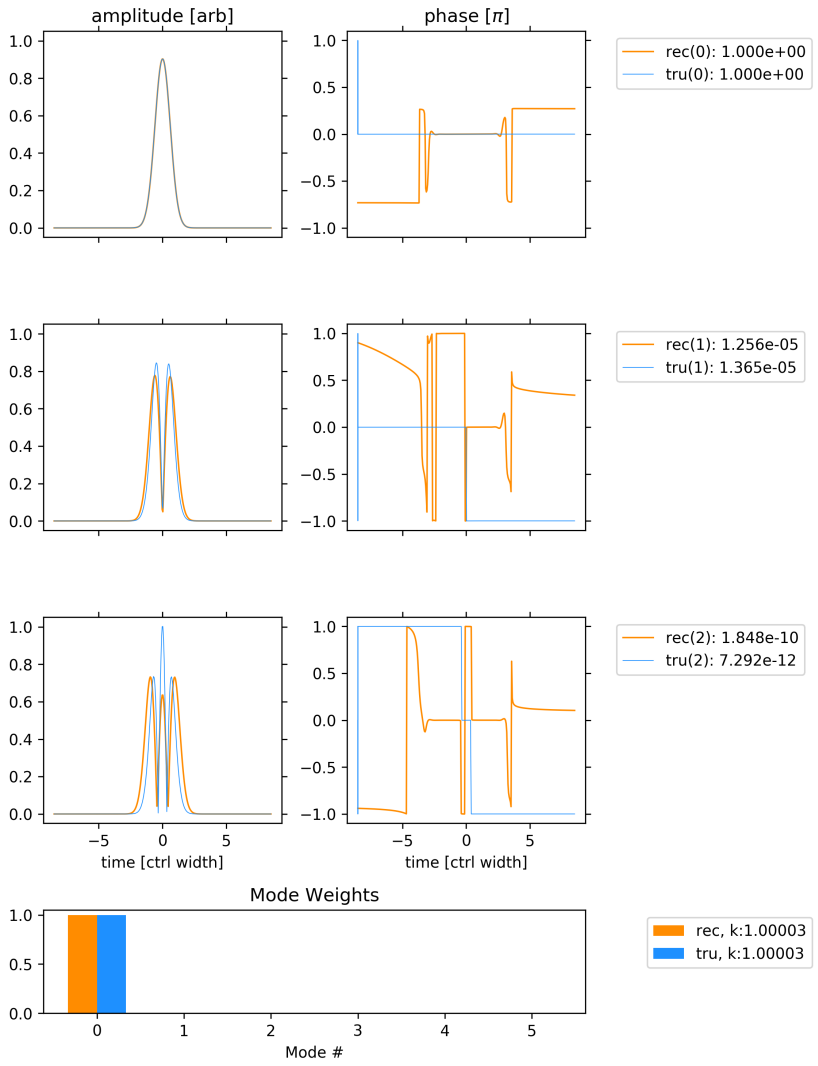
In this case we see a much more favourable agreement between reconstructed and simulated modes and their weightings, for the first two eigenvalues. Of course we anticipate this to begin to fail at $d = 3$ where the temporal mode basis of the probe is truncated.

Comparing these two figures, we note that the reconstruction also captures the mode distortion seen in the simulation: that is temporal shift of the peak intensity[18] time of the eigenmodes, and asymetric distortion in both the amplitude and phase, as the power is increased.

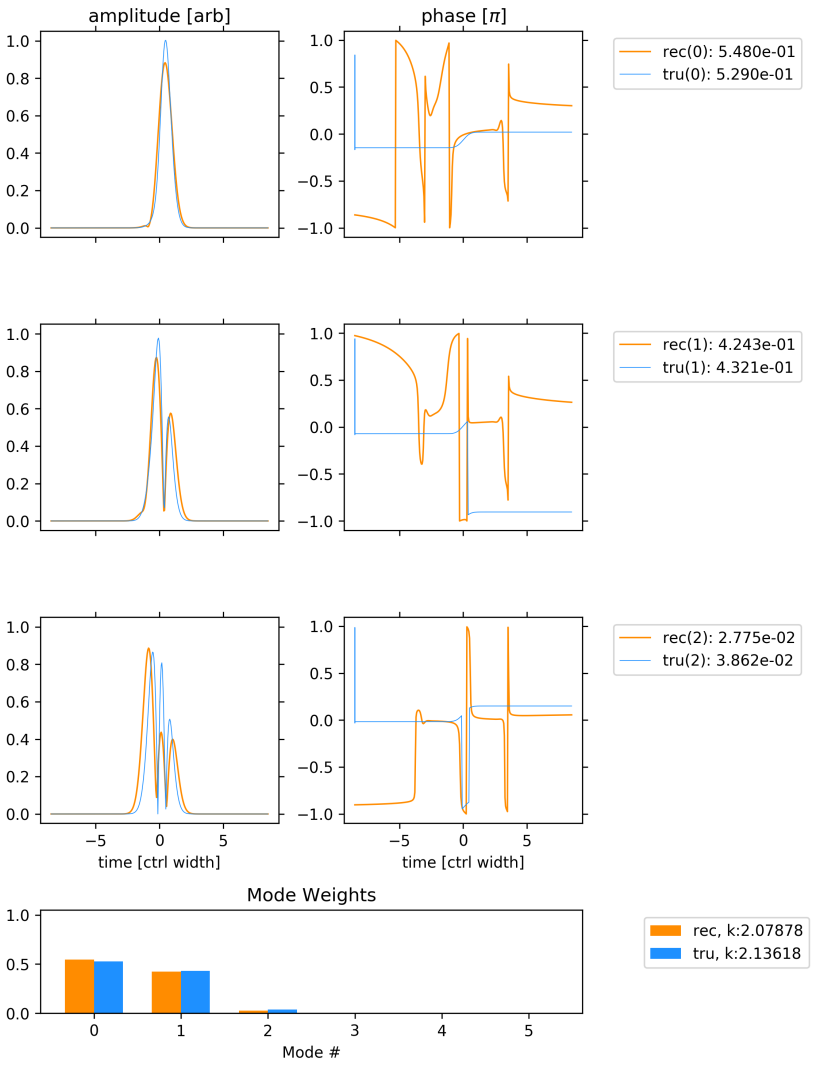
As expected, the agreement between reconstructed and simulated modes deteriorates as the effective mode number of the interaction increases.

The last case we consider is again with ideal Hermite-Gaussian states and complementary MUB measurements at low power (in the effectively single-mode regime, yielding $k = 1.0003$ for the underlying state), but which are either (a) for zero-delay states in 3D mismatched with the control (12 measurements), or (b) for time-scanned states in 2D which are matched – at 0 and ± 1 ns (18

[16] More explicitly: the $|1\rangle$ efficiency is ~ 0 in both cases, whilst there is a factor of $\sim 3\times$ between the $|0\rangle$ efficiencies.
[17] With the simplified model of the memory interaction.
[18] Although the "centre–of–mass" time of the pulses remains stationary.

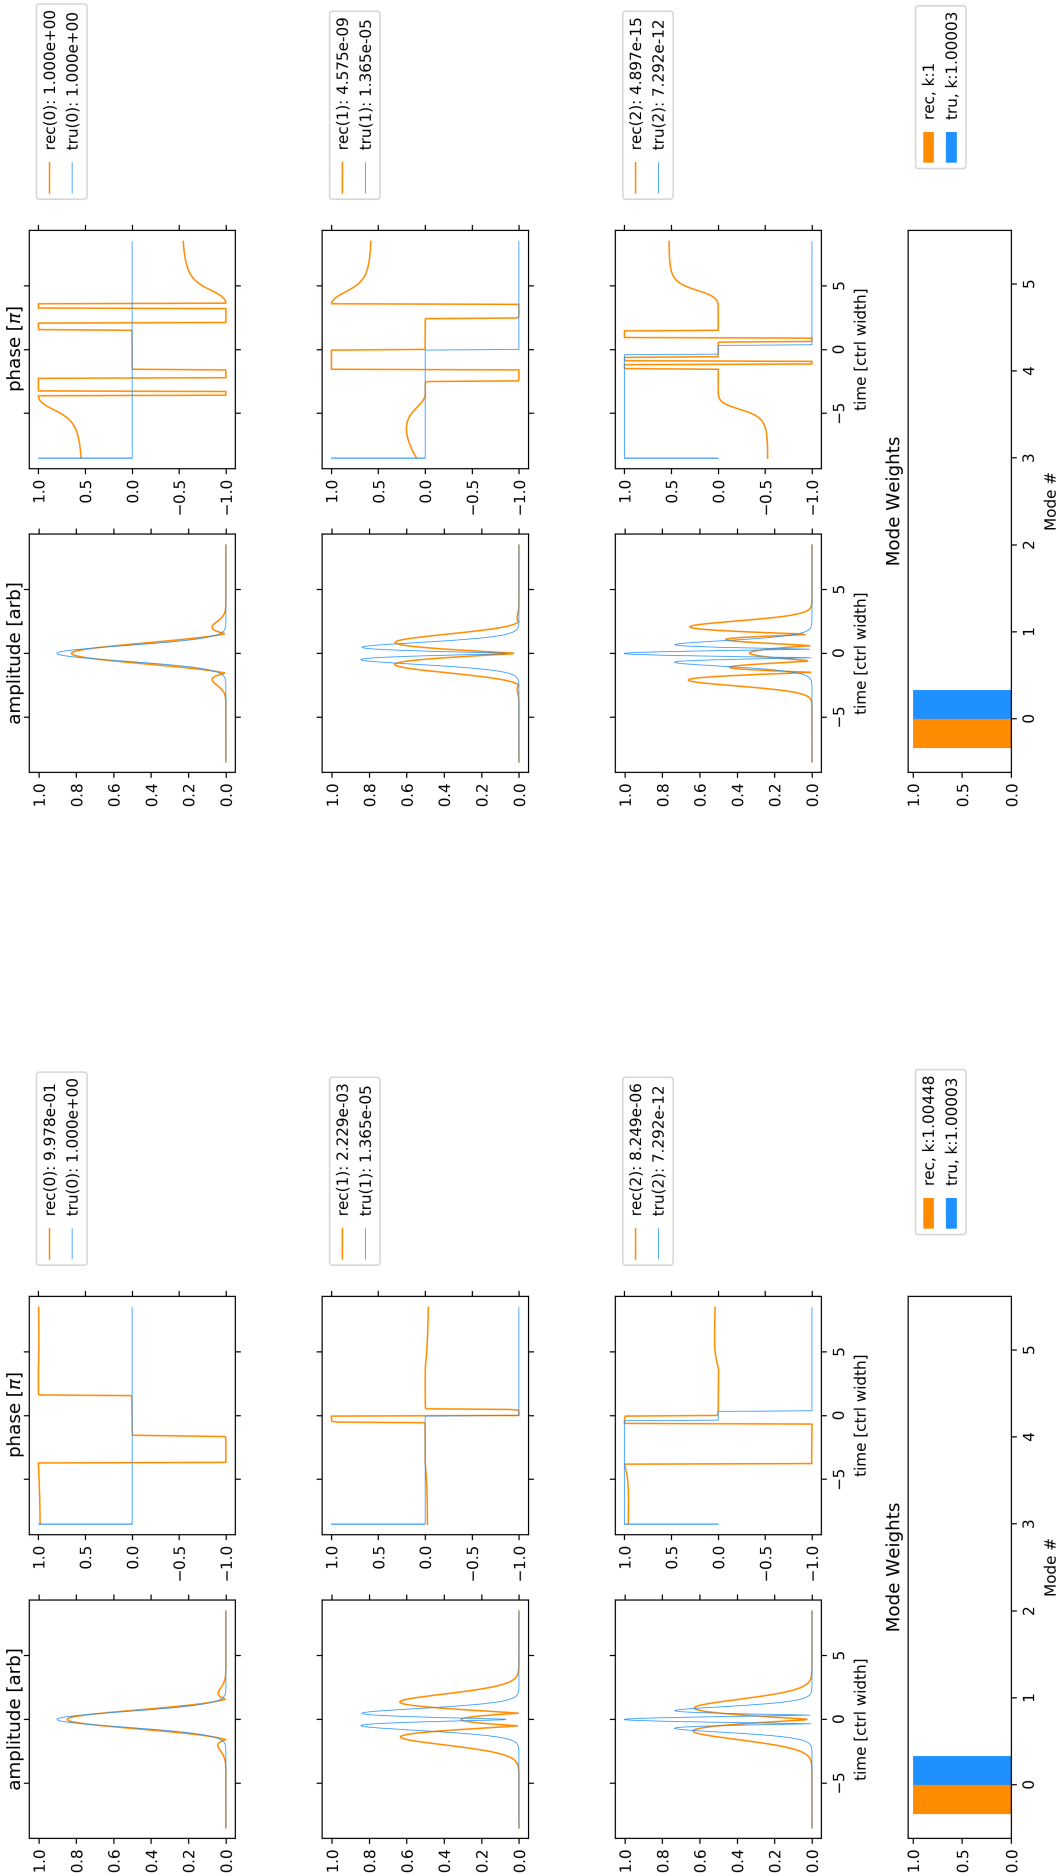


(a) Low power (W_0)



(b) High power ($200W_0$)

Figure 7.10: Reconstruction using simulated efficiencies for width matched probe states (ideal MUBs) in 3D with 12 MUB states at zero delay. (a)-(b) simulated at different powers. These display excellent agreement between the eigenmodes and mode weightings of the reconstructed and underlying Green's function, particularly at lower powers (and hence lower effective mode numbers).



(a) $3d$ probe basis, with no time information, width mismatched

(b) $2d$ probe basis, including time information, width matched

Figure 7.11: Reconstruction using simulated efficiencies with (a) 3D MUB probes 50% longer in duration than the control, (b) 2D timescanned MUBs, width matched.

measurements). The dimensionality here nD refers to the order n at which the HG probe basis is truncated, and defines the subspace about which we can extract information.

The width-matched 2D states here seek to give insight into the "best case" information that we can extract from time-scanned probe states to attempt to dilate the space supported by the probe states.

Qualitatively, we see similar results between the two cases, where the algorithm returns an appropriate width of the principle eigenmode, although with some lingering artifacts – seen in the side-lobes, and the widths of the higher order modes are poorly matched.

This indicates that in these measurement configurations we can only retrieve limited information about the underlying ρ, and the measurements are, in some sense, not optimal. Furthermore, in both cases we fail to extract information about the higher order modes, although due to the low mode number in this case, the resulting "structure" of the higher order eigenmodes is not physically meaningful.

7.5.3 Conclusions

We note in the appendix D, that for an informationally incomplete measurement, in general maximising the likelihood cannot uniquely define the state; and instead application of the algorithm will result in one of a family of states.

The structure of the modes reconstructed here point towards this lack of information, and so by adapting our measurement choices in simulation we aim to make use of "more insightful" measurements i.e. ones that provide more appropriate constraints with respect to which we maximise the likelihood.

7.5.3(a) Memory Reconstruction

Restricted to a low dimensional space, the algorithm is particularly sensitive to the choice of measurement states. As the efficiency of the memory increases a higher dimensional vector space is required to describe – and therefore access information – about the interaction matrix, since (i) the effective mode number of the interaction increases, and (ii) the distortion of the eigenmode basis themselves from the unperturbed Hermite–Gaussian states, means that decomposing the interaction eigenmodes into a *distinct* orthonormal basis necessarily requires a higher dimensional vector space in the unperturbed coordinate basis.

We proceed further with the width-matched 3D zero delay and observe the performance of the reconstruction as W is increased, and in figure 7.12 we observe the eigenmode basis of the memory evolve. Whilst at low coupling, the principle eigenmode closely resembles the control profile, as the coupling is increased we observe two characteristic features of (i) the leading peak of each mode increasing in height, and (ii) the shift in peak time. Qualitatively there is close agreement between the eigenmodes of the reconstructed state, and the original simulated state, up until the dimension $d = 3$ where the probe space is truncated, at which point we can no longer distinguish the contributions from higher order modes.

The distortions in amplitude seen here (and correspondingly the modified phase profile we observed in figure 7.10) are directly analogous to the skews seen in the Schmidt modes of the quantum pulse gate as the coupling is increased [211]. This characteristic distortion of the modes (along with driving the process to become more multimode) as the coupling is increased is a general physical property of

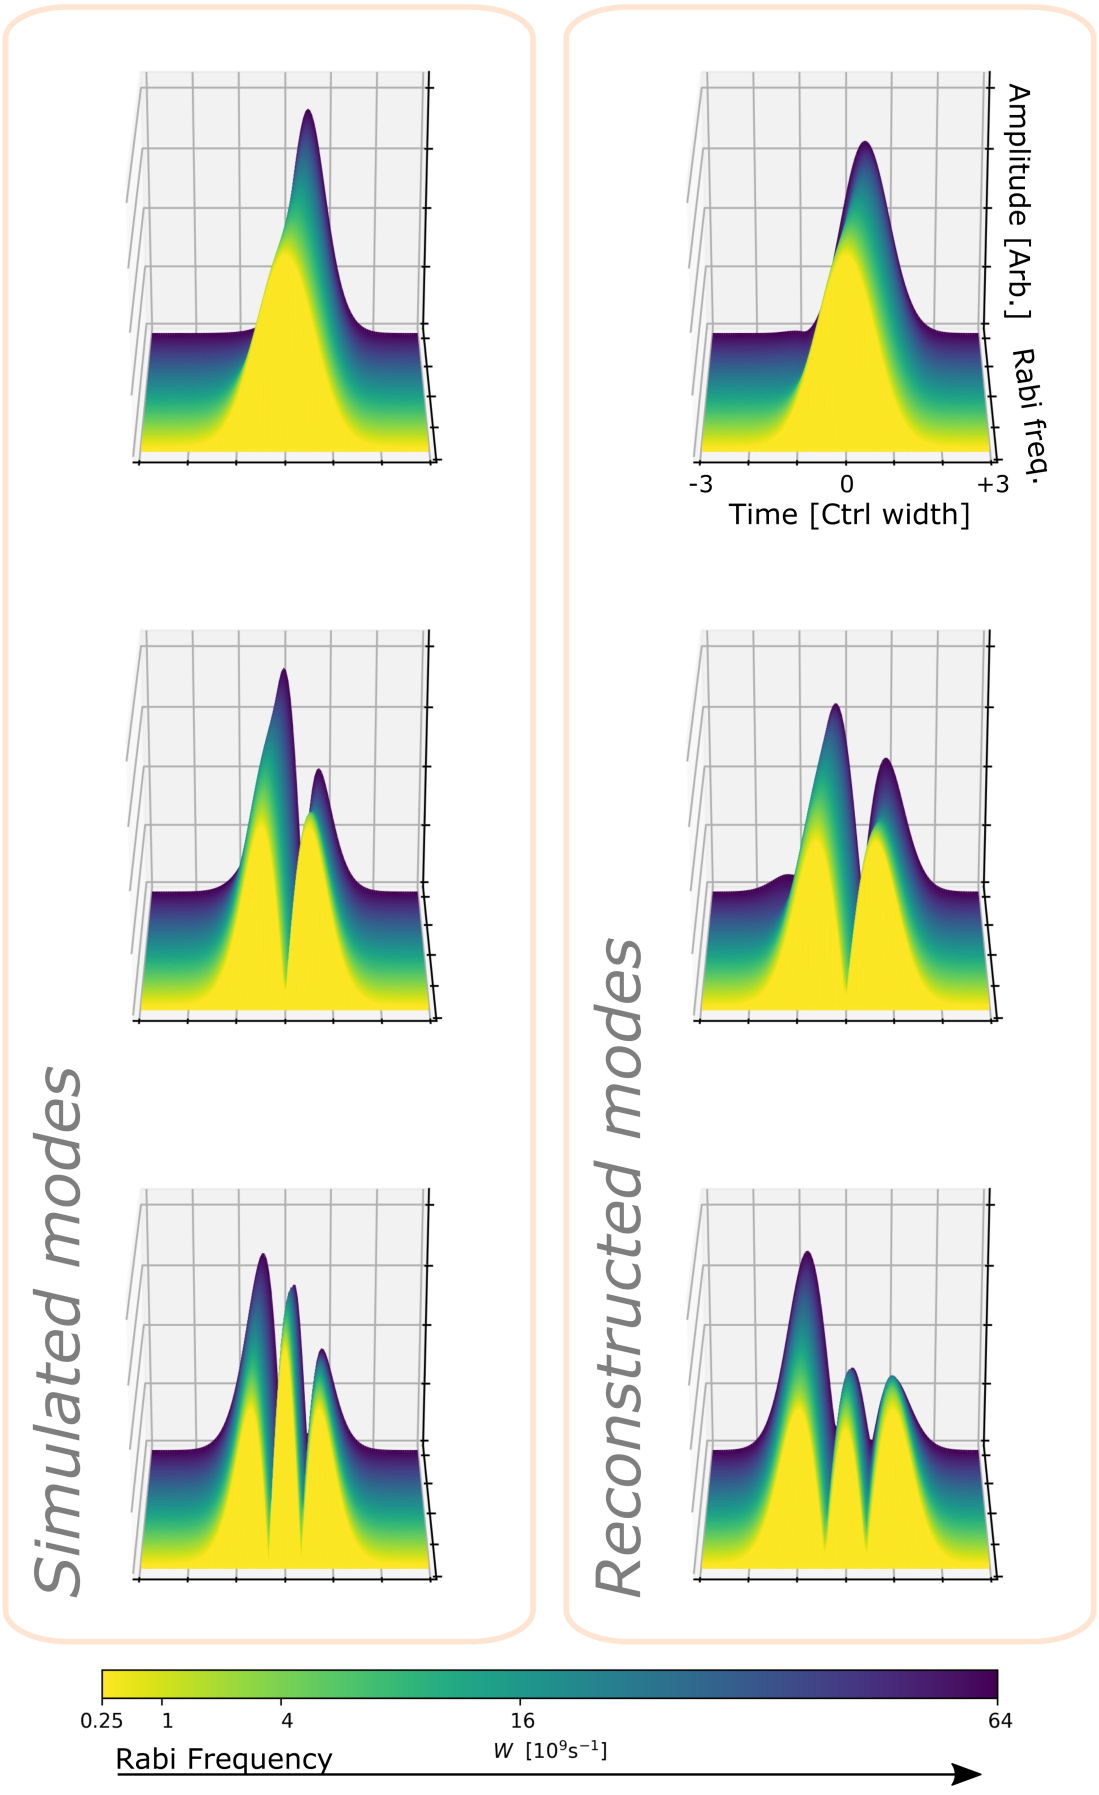


Figure 7.12: Intensity temporal profile of the simulated and reconstructed (from simulated data) interaction eigenmodes as the coupling parameter W is increased. This displays a qualitatively close match between directly evaluated and reconstructed eigenmodes.

the "broadband beamsplitter" at elevated coupling, due to time ordering effects [212–214].

The origin of this effect may be seen by considering the coupled equations of motion for the relevant fields, for instance from equations 2.50:

$$\partial_x \begin{pmatrix} S_z \\ B_\tau \end{pmatrix} = \begin{pmatrix} \frac{-\chi^2}{c\Gamma} & \frac{\mathrm{i}\chi\Omega}{\Gamma} \\ \frac{-\mathrm{i}\chi\Omega^*}{\Gamma} & \frac{-|\Omega|^2}{\Gamma} \end{pmatrix} \begin{pmatrix} S_z \\ B_\tau \end{pmatrix} = \boldsymbol{M}(\tau) \begin{pmatrix} S_z \\ B_\tau \end{pmatrix}, \tag{7.13}$$

where ∂_x is the derivative over the appropriate degree of freedom, x. Since the control field temporal mode $\propto \Omega(\tau)$ is time dependent, the matrix governing the coupling in general does not commute at unequal times: $[\boldsymbol{M}(\tau), \boldsymbol{M}(\tau')] \neq 0$.

Whilst at low coupling strengths, the conversion mapping signal S to spin wave B is small – this occurs uniformly across the $z-$distribution of atoms ($B(z) \approx$ const.). In this case, the converted field is weak, and the interaction depends solely on the overlap between the signal and control temporal modes, and is hence single mode. At elevated coupling, however, the conversion no longer occurs homogeneously across the cell: there is initially a strong conversion at the start of the cell,[19] at which point the interaction now depends on a three-mode overlap – and back-action between the storage mode and signal occurs. This may be interpreted as simultaneous read-in and read-out, and results in a distortion of the eigenmodes, as well as driving the interaction to become multi-mode.

In order to describe this, the time evolution of the fields must be handled appropriately. The time-ordering analogue may be seen with the unitary evolution of a time-dependent Hamiltonian:

$$\hat{\mathcal{U}}(t) = \mathcal{T} \exp\left[-\frac{\mathrm{i}}{\hbar}\int_{-\infty}^{t} \mathrm{d}t' \hat{H}(t')\right] = \exp\left[\sum_j \Omega_j(t)\right], \tag{7.14}$$

where $\mathcal{T}$ imposes time-ordering operation – which, for instance, may be computed using a Magnus expansion (with a form as in the second equality above) [215]. This operation is necessary if $\left[\hat{H}(t), \hat{H}(t')\right] \neq 0$, which is equivalent to the non-commuting matrices $\boldsymbol{M}(\tau_0)$ and $\boldsymbol{M}(\tau_1)$ above [214].

Lastly to demonstrate the limits to the information we may extract with this approach, figure 7.13 shows the mode weightings, again as a function of W (where the colour scale is drawn from the simulated *optimal* efficiency). We observe that as the coupling increases and the higher order modes pick up, we have good agreement with the simulation (grey) in the single mode regime, and discrepancies become apparent as the coupling is increased.

This is clarified in figure 7.14 in which we simply compare the extracted mode number as a function of W and see the reconstruction begin to noticeably fail around $k \sim 1.5 \sim d/2$, which occurs at $W \sim 23 \times 10^9$ s^{-1} in this model, at these operating conditions. This means that this method would be sufficient to verify the *single modedness* of the interaction for the HG0 control field pulse.

We note that the (simplified) Green's function used to model these data is independent of the *shape* of the control pulse, for a given pulse energy [139]. More explicitly, the Green's function is defined by the normalised time coordinate – and a simple coordinate transformation can be constructed for a given temporal profile in order to obtain the mapping between ω and t. This has the consequence that *if the Green's function is verified to be single mode for one control pulse, it is necessarily single mode for all other pulses containing the same pulse energy.*

[19] ...which then decays with z. In the model used here, $B(z)$ has the form of a Bessel function (refer to equation 7.1).

Of course, this is a simplified model which does not capture the full physics.[20] In any case, in our experimental operating regime, with low coupling, we expect this to be a good model,[21] and thus, this provides a rather promising basis on which to move the experiment forward.

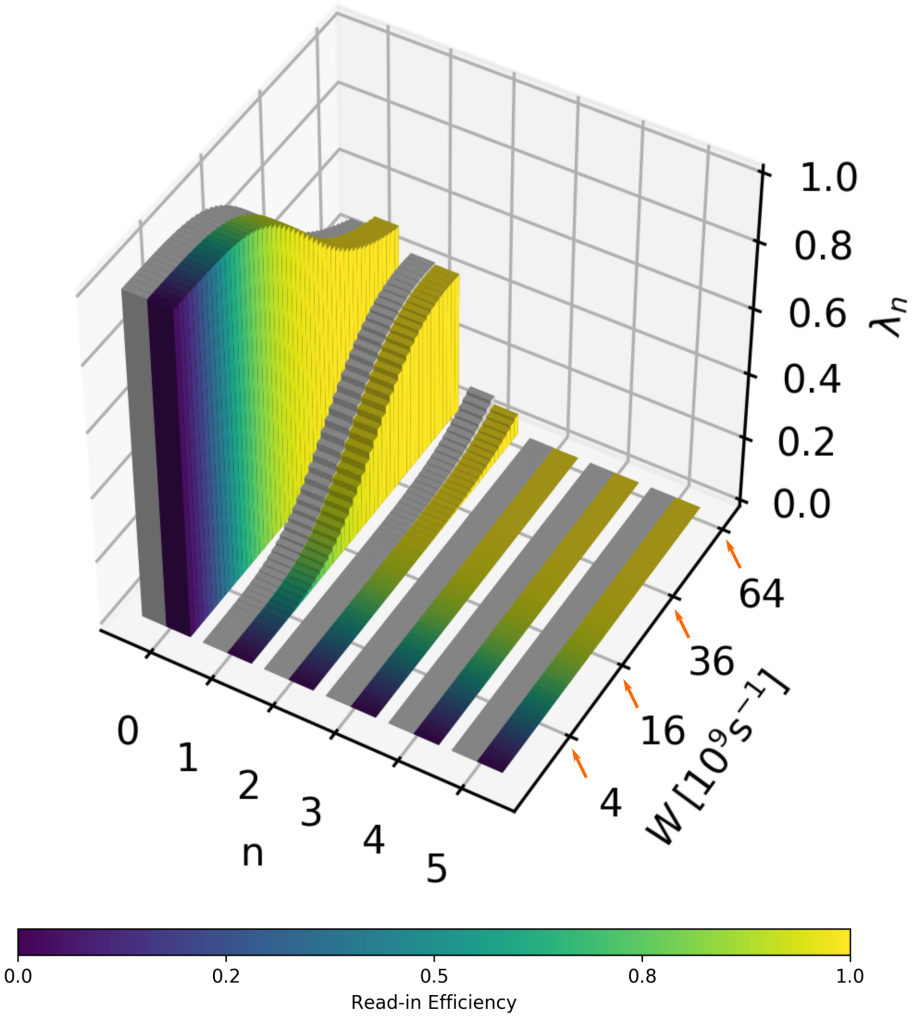


Figure 7.13: Relative weightings of the first 5 eigenmodes from simulated (grey) and reconstructed (colour) matrices. The colour scale indicates the modelled *storage efficiency* of the principle eigenvalue.

7.5.3(b) Consequences for Mode Selectivity

These initial results are a promising demonstration of exploiting the single mode property of the Raman memory. We have demonstrated the selective storage of the Raman interaction, and the suppression of the storage of a (near) field orthogonal state. With additional measurements in the mutually unbiased bases we are able to *begin* to extract information about the temporal mode structure of the memory interaction. Including some extra information from temporally delaying these probe states we are able to recover more information than otherwise available; although, further simulations based upon these results suggest[22] that more information will be available by moving to higher

[20] Even with FWM the $t \leftrightarrow \omega$ coordinate transformation can still be used, and the Green's function depends only on the pulse energy.

[21] Based upon (a) the agreement for the lower power data with the model in figure 7.7, and (b) our estimated operating regime ($\Omega_{\max} \sim 1 \times 10^9$ s^{-1} c.f. $\Delta = 2\pi \times 15.2$ GHz).

[22] And perhaps intuitively so.

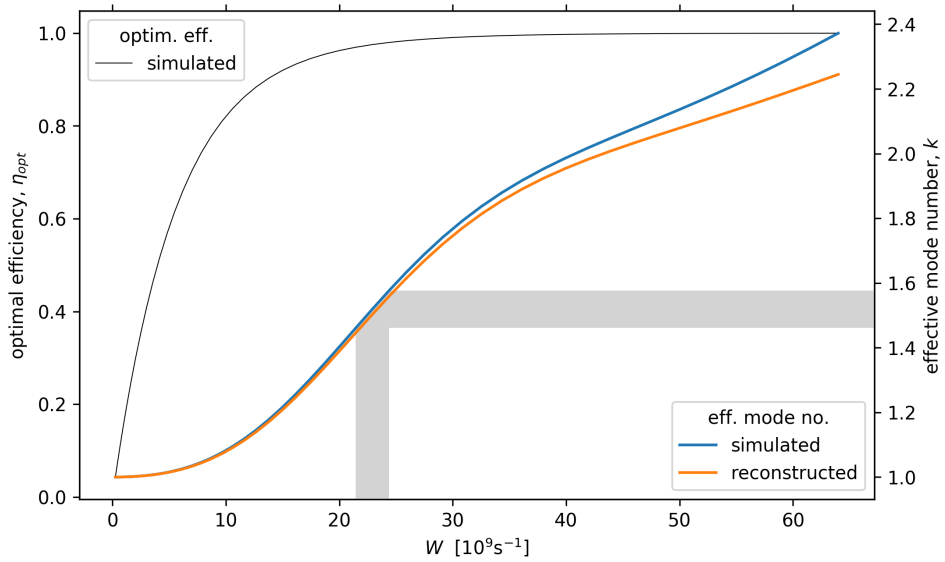


Figure 7.14: Storage efficiency of the optimal mode for the simulated Green's function as a function of W. The effective mode number directly/reconstructed from the simulated Green's function is shown on the right hand scale, which become noteably different around $k = 1.5 \sim d/2$.

dimensional probe states, particularly if the probe and control pulses are more closely matched in duration.

In fact, we show that (width-matched) $d = 3$ probe states should be sufficient to verify the single modeness of the interaction – directly accessible in experiment for a control pulse addressing HG0 and HG1. In this case, uncertainty in probe preparation fidelity will likely be the dominant error, and the means of directly characterising the optical fields (using methods in appendix B) will be valuable.[23]

For practical applications in QIP in general scaling up the dimensionality available offers enhancements in communication protocols. For example in QKD offering enhancement in the secret key-rate as well as increasing tolerance to error in the channel. However, this brings with it additional overheads in the practical requirements to generate, manipulate and interface with an expanded Hilbert space. A recent experimental study of different QKD protocols using OAM states [216] suggests that in this case $d \sim 4-5$ might in general be optimal since whilst increasing the dimensionality increases the security and information extracted per *sifted* photon, the fraction discarded (due to mismatched measurement basis) increases, and additionally working with a higher dimensional encoding increases experimental complications.

In terms of building on these results, and working within a restricted space, this method seems promising to verify the mode structure of the Raman interaction for the low dimension inputs.

The present results were obtained at memory efficiencies between $20-50\%$ and as observed in the numerical investigation, as with analogous work on the QPG, it is a general physical property of the broadband beamsplitter that as the efficiency of the process is increased, the interaction becomes more multimode, and the basis of the interaction is rotated.

Refering back to the previous part of this thesis, we note that in principle the cavity Raman memory, with sufficient engineering, would offer the prospect of attaining high total efficiencies, whilst

[23] This provides an approach to evaluate the uncertainty in the reconstructed Green's function, via a Monte-Carlo method repeating the algorithm with perturbed probe states with physically reasonable errors.

remaining a purely single mode interaction.

Analogous to the boundary conditions imposed by the cavity, a proposed approach to circumvent the increase in mode number at elevated coupling in the QPG is to concatenate two crystals at 50% conversion efficiency [217, 218], although this necessitates temporally shifting the pulses back to their initial relative positions [219].

Nevertheless, single modedness aside, the mode distortion of the memory eigenmodes as the efficiency is increased would remain a problem. That is, that the mode addressed by (and that – not necessarily the same state – retrieved from) the memory become increasingly skewed, in amplitude and phase, with respect to the unperturbed TM alphabet encoding. As illustrated[24] in figure 6.1 in the cavity case, as well as more generally, we can consider inverting the problem,[25] and include shaping of the control field to address a given target mode such that we can selectively address an arbitrary temporal mode.

7.6 Summary

The methods here, along with TEA+PARTI (see appendix B), add to the toolbox for interfacing with temporal modes of light, and with these results are the first steps towards employing the Raman memory as a device for the manipulation of photonic temporal modes.

The results demonstrate a high degree of separability of the Raman memory, $\mathbb{s} \approx 98\%$, at read-in efficiency of $\eta_{\text{in}} \approx 20\%$. We have identified a method for verifying the structure of the interaction, and its single modeness, via the reconstruction of the Green's function from memory measurements; and show that, extending to a 3D input state basis, closely matched with the control will be sufficient to reliably benchmark the selectivity in the low coupling regime.

Building upon these promising results, the experiment is moving forward. The first extension is to allow for arbitrary shaping of the control field. This will enable us firstly to verify a full process tomography of the memory interaction (extending from operations solely on states as defined by the Ti:Sa derived control pulses[26]) in that it operates on any letter – and upon rotated bases – in the TM alphabet.

Promoting the Raman memory as a device for manipulating temporal modes in situ, the control infrastructure developed for the AWGs will be readily applied to demonstrations of mode conversion[27] for instance bandwidth conversion. This would be an enabler for building up large scale quantum networks, and interfacing different elements since the ideal spectral operating parameters of these components are not in general compatible – where in contrast to a passive device such as a filter which may only narrow the bandwidth, at the expense of transmisson, the memory interaction allows for both expansion and contraction on read-out. The upper bandwidth here would be physically limited by the ground state splitting[28] (9.2 GHz, in Cs) but also technically may be limited by achieving the same pulse energy in a shorter carved control pulse; the lower limit on bandwidth is likely to be set by the atomic dynamics (∼10 MHz).

[24] Note that this does not arise from the time-ordering effect discussed, since in the cavity case we may remain in the low coupling regime, enhanced over many passes. This distortion is a manifestation of the a.c. Stark shift.

[25] Or experimentally optimising upon the shape of the control field.

[26] If both signal and control are derived from the same laser, then the relative frequency drift between the two fields would no longer be an issue.

[27] c.f. retrieved mode from QPG defined solely by mode matching.

[28] Bearing in mind we should also operate in the adiabatic limit.

7.6.0(a) Key Results

This chapter realises the first step towards using the Raman memory interaction for interfacing with temporal modes of light, and in so doing developing a toolbox of methods for future experiments.

These results first experimentally demonstrate the application of using the Raman memory as a temporal-mode selective device; and secondly, to verify the mode-selective performance of the device based upon reconstructing the Green's function for the interaction.

§7.2 I introduce the conceptual methods for tomographic reconstruction and the models used for numerically describing the memory interaction.

§7.3 I outline the experimental details including the state preparation, and the changes to the memory infrastructure from the first part of this thesis.

§7.4 I describe how one can reconstruct the Green's function from measurement of the memory efficiency. This draws upon methods in quantum state tomography using informationally incomplete measurements, and I then discuss how this can be mapped to our experiment.

§7.5 The results of applying temporal-mode tomography to the Raman memory from experimental data are presented. I also supplement this with further simulation investigating how, with experimental imperfection, the choice of measurement may be used to improve the performance in future.

CHAPTER 8

CONCLUSIONS PART II

8.1 Summary

I have presented my initial investigations into exploiting the temporal mode selectivity of the Raman memory interaction and numerically explored measurements for verifying the modal structure of the memory. I developed the infrastructure in our lab to facilitate interfacing carved signal pulses with the Raman memory, including the python libraries for running the experiment.

Whilst not used directly in these results, I also direct the reader to appendix B in which I present an additional method for this toolbox which enables the direct reconstruction of the complex spectrum of pulses with $\gtrsim$ ns duration, suitable for characterising temporal modes (and testing the pulse-carving infrastructure) compatible with quantum memories. We anticipate that this would complement the methods presented here for future evolution of this work, with the cavity memory targeted at slightly more narrowband signals

8.1.1 Mode Selectivity

We experimentally probed the structure of the memory interaction by measuring the memory efficiency with which different coherent state temporal modes were stored for a two-dimensional basis of the lowest order Hermite-Gaussian functions, and linear combinations to construct near mutually unbiased bases. The storage efficiencies for these six states were recorded as a function of time delay at a range of control field powers

At low coupling, we expect the mode addressed by the memory to be similar to the control pulse temporal mode. In this regime, we can estimate a separability $\mathbb{s} \approx 98\%$ from the experimental reconstruction in figure 7.9 (a). We do not place too much weight on this number itself, acknowledging the limited information available to the reconstruction method with the current experimental infrastructure as compared with the performance of the "best case" 2D time-scan performance in simulation. However this does provide a promising foundation for the next generation of work realising the Raman memory platform for the manipulation of temporal modes.

We have shown, through simulation, that further improvements to the results here are available if (i) the probe state basis and the control field are more closely temporally matched, and (ii) if we could extend to a higher dimension of probe states. In parallel to this, based upon these results a clear immediate step is to also shape the control and demonstrate the same strong suppression when we use the memory to address different temporal modes.

8.1.2 Mode Reconstructions

Through numerical investigation we show that with imperfect state preparation (and thus suboptimal spanning of probe vector space) and when the dimension of the probe vector space is low (and isn't sufficient to *support* the control field or interaction matrix) this limits the information accessible for the reconstruction of the interaction via measurement of the memory efficiency. We can partially compensate for this by incorporating time scan information of the probe states, but with access to the first three Hermite-Gaussian basis states, and associated mutually unbiased construction the tomography algorithm performed excellently in reconstructing the Green's function in the weak coupling regime.

Further information could be gained with access to even higher dimension of HG states ($d > 3$), if one can choose alternative linear combinations of the "computational" basis states, if they allow for smoother electrode signals. Analogously to the time scans, these linear combinations need not necessarily be mutually-unbiased, in which case they will not be optimal in an informational sense, but still would be sufficient to span the necessary vector space if this allows for more faithful preparation of probe states.

The immediate next step from the results presented here here would be to extend to reconstruction of the interaction for different control profiles – that is *experimentally* showing mode selectivity for the storage of different states.

8.2 Future

The results presented here show the first investigation, and demonstration, of mode selectivity in a Raman memory.

Many of these issues faced here can be addressed if the signal and control are derived from the same laser: whereby both signal and control pulses are carved from the same diode laser then using a tapered amplifier (TA) to amplify the control pulse to sufficient intensities. As a first consequence, the improved temporal match between the signal and control would enable a moderate increase in efficiency simply due to the favourable overlap of the pulses.

Furthermore, with a carved control, as opposed to a pulse from the Ti:Sa, we would be able to work with temporally longer pulses, and prepare states with better fidelity (suffering less ringing of amplifiers) up to higher temporal mode dimensions – and as discussed in the previous chapter, this would improve the information which may be retrieved via reconstruction of the memory Green's function from measurement of the storage efficiencies, such that one can robustly experimentally verify the selectivity of the interaction.

The extension of this experiment to a width matched 3D probe states at zero delay is a realistic near term goal. Based upon the reconstruction method applied to simulation in this case, which performed robustly with excellent agreement at lower coupling (and hence low mode numbers), this should provide a confident basis on which to evaluate the separability and single mode performance of the memory in practice.

In practice, this change would also improve the experimental stability by removing the relative frequency drift between the signal and control, as well as opening the door to accessing and *interacting with* a higher dimensional temporal mode basis.

Also in the scope of improved experimental stability comes addressing the observed FMZI bias drift over the course of an experiment. This could be achieved with a commercial solution, or employing

some of the methods in the appendices, for instance interleaving a protocol to minimise the CW transmission through the FMZI, in between carving experiments, and employing an adaptive algorithm to find the best electrode parameters (gain, crossing point and signal) for each state to be prepared.

Refering back to Part I of this thesis, recall also, that with freedom to carve the control field we may be able to compensate the a.c. Stark shift and enable higher efficiencies to be reached. With an operational TA, this would also be a step towards a "quantum" implementation, whereby we could use an arbitrarily carved pulse to pump a PDC source as a source of photons in given temporal modes.

Additionally, with longer pulses, these would be compatible with our current TEA+PARTI set up,[1] and we could obtain a direct optical characterisaton of the signal and control pulses which would provide a more rigourous basis for obtaining the projective elements used in the reconstruction.

The results presented here thus provide the groundwork, and establish some of the tools, for a wide range of applications exploiting the Raman memory as a device for the manipulation of temporal modes of light. I discuss some appealing near term applications in the following sections.

8.2.1 "Stark Compensated Cavity–Magic–Raman" Memory

Whilst we demonstrate in this first investigation that, with the bulk Raman memory, we may reach a high degree of selectivity, this is necessarily accompanied with rather modest efficiencies $\sim 20\%$. Combating the Stark shift will enable higher efficiencies to be reached for a given selectivity (see figure 7.6), although in free space this increase will necessarily be small in the regime in which the memory interaction remains approximately single mode.

To overcome this limit, we recall the conclusions of part I of this thesis. With sufficient engineering, introducing a low loss optical cavity around the memory can simultaneously enable high efficiencies to be reached whilst remaining single mode. An $M = 0$ cavity, with control pulse shaping offers the most enhancement in terms of efficiency and noise suppression;[2] which would constitute a scalable platform for the local manipulation of temporal modes.

The nearer term design, simulated in figure 4.10 should simultaneously enable efficiencies $\sim 50\%$, lifetimes of ~ 1 μs, and be sufficiently low noise to demonstrate single mode storage of ~ 100 MHz bandwidth, quantum signals. One could also envision implementing the "Magic Raman" protocol (introduced in section 5.3.2) to provide further noise suppression.

Furthermore, we note that the control profile independence of the Green's function *for a given pulse energy* (as was the case for the model used in the previous chapter) is also a feature of the theory for the cavity-Raman interaction. In principle this means that once the Green's function has been extracted for some pulse shape $\omega(t)$, this may be related by a coordinate transformation $\omega(t) \leftrightarrow t$ to any arbitrary control pulse ω'. Investigating this will give an independent metric of the validity of the models used.

With improvements to the pulse shaping infrastructure to carve the necessary control pulses, together with insights into the mode structure of the memory interaction, one should be able to selectively store (and retrieve) arbitrary target basis states (a) without distortion at high efficiencies, and (b) to reach high efficiencies by compensating for the Stark shift roll-off. This would pave the way for a system that would be suitable for quantum implementation – that is, with increased stability,

[1] Although, some further experimental characterisation is required.

[2] Note that, in order to achieve sufficiently low loss and avoid the "turn over" in figure 4.9 , the buffer gas pressure may be reduced in order to minimise pressure broadening leading to linear absorption at the signal frequency – thus limiting the lifetimes that may be reached.

lower loss, and reduced noise – the system could be readily interfaced with single photon temporal modes.

8.2.2 Photon Source Purification

If one can combine these to realise a high efficiency, low noise memory suitable to be interfaced with single photon TMs, one immediate application would be to exploit single-modedness of the cavity memory as a "filter" which can be used to purify a photon source.

A source of pure photons is a vital resource for photonic QIP – where we require photons to be indistinguishable and pure in order to observe two-photon interference [220]. Realistic photon sources may emit into many modes, and – for instance in PDC – the detection of a herald, projects the remaining photon into a mixed state. To circumvent this usually requires engineering of the source e.g. [221, 222], which presents a new challenge for the generation of pure photons for each individual application – for instance at some specific wavelength.

A simple method to address this is to resort to spectral filtering which enables the heralding of approximately pure single photons; however this necessarily would detriment the brightness of the source [223].

A single-mode TM interface has been proposed [184, 194] as a versatile solution. Using the broadband beam-splitter to address a *single* mode from the herald arm and detecting the retrieved signal will herald a photon in a pure state, corresponding to the herald mode addressed by the TM selective device. This process will, of course, impact upon the achievable brightness of the source, as a consequence of finite device efficiency, neverless with this infrastructure we may herald a pure photon in a desired temporal mode as opposed to a spectrally filtered photon yielded by passive filtering of the herald.

8.2.3 Mode Conversion

Another appealing feature of using the memory to manipulate TMs is that it enables the in situ conversion between encodings. A simple example, already mentioned, is between TMs and time bins, where one would need to selectively delay components of a signal and thereby map each TM to a different time bin.

The memory also enables bandwidth conversion between read-in and read-out [119]. Again, if one compares this with a passive filter which cannot increase bandwidth, and will introduce further loss the more the signal is filtered out.

Reshaping between read-in and read-out(as demonstrated in EIT in [101]) also makes the single mode memory an ideal system to interface between dissimilar systems – for example mapping a signal from the emission profile of some matter node, into a different optical mode suited another part of a network. In particular, one might wish to exploit temporal mode encoding for transmission down single mode fibre, and then interface with another (potentially different) matter system at the other end.

These material systems will in general operate at dissimilar wavelengths, bandwidths and we may desire different amplitude/phase temporal profiles for optimal interaction. This might be different species of emitter, or perhaps more promisingly, emitters in the *same* solid state system, but each subject to its own local environment giving rise to perturbed spectral properties.

The Raman memory therefore constitutes a versatile platform which enables interfacing these

different systems, offering freedom to translate between different bandwidths and centre frequencies over the MHz – GHz regime.

8.3 Outlook

In this thesis, we have identified – and sought to address – two elements of the Raman memory protocol in warm caesium vapour.

Firstly, we introduced the cavity Raman memory, in which we demonstrated a reduction of noise compared to previous implementations. Then we presented a theoretical and numerical investigation of how to optimise this system on a number of fronts. These efforts work towards facilitating the design (and future implementation!) of a low noise, integrated system with low experimental overheads; which could constitute a valuable addition to the toolbox for photonics, within a wider quantum information processing architecture.

Then, having identified the feature of temporal mode selectivity of the Raman memory, we have performed the first experimental investigation of the mode structure of the Raman memory interaction. This provides the basis for using the Raman memory as an interface with TMs, which have been identified as an appealing architecture for photonic QIP. This platform holds promise to be a versatile device for interacting with temporal modes, and will enable a new generation of experiments and operating regimes exploiting the high-dimensional robust encoding offered by temporal modes.

8.4 Final Thoughts

On paper, the Raman memory in warm alkali vapour occupies the middle ground in the quantum memories landscape: in comparison with the state of the art (see table 1.1), it does not yet offer the 87% efficiency of GEM in [121], the 40 s lifetime of EIT in [130] nor the THz bandwidths of Raman in diamond in [131].

However, a non-classical autocorrelation has yet to be demonstrated as being preserved in any warm vapour $\Lambda-$based memory interaction: GEM, EIT or Raman. In the first part of this work, we have seen how this could be realised in a compact package with bandwidths still exceeding those of many other protocols. This still remains a technical challenge, nevertheless, the table-top demonstration of successfully suppressing FWM noise – together with the optimisation approaches – presented here make for a promising start.

From there, we have also begun to explore a new application of the Raman memory, based upon the single-mode nature of the interaction, as a device for manipulating temporal modes of light. This sits well in the need for the controllability of quantum systems. Again, this is just the start. But so far we have laid some of the groundwork and established a toolbox to facilitate making use of the Raman memory as a device for interfacing different systems and encodings.

If these two themes are united, then I think this platform – or at least what we have learned from it – could find a place as a useful enabler, in whatever directions the field of quantum information science moves in the future.

BACK MATTER

BIBLIOGRAPHY

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

[2] V. Giovannetti, *Quantum-Enhanced Measurements: Beating the Standard Quantum Limit*, Science **306**, 1330 (2004), 0412078v1.

[3] H. J. Kimble, *The quantum internet*, Nature **453**, 1023 (2008), 0806.4195.

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

[5] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Quantum machine learning*, Nature **549**, 195 (2017), 1611.09347.

[6] 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).

[7] D. Deutsch and R. Jozsa, *Rapid Solution of Problems by Quantum Computation*, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **439**, 553 (1992).

[8] L. K. Grover, *A fast quantum mechanical algorithm for database search*, Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96 , 212 (1996).

[9] P. W. Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, SIAM Journal on Computing **26**, 1484 (1997), 9508027.

[10] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, *Quantum algorithms revisited*, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **454**, 339 (1998), 9708016.

[11] R. Raussendorf and H. J. Briegel, *A One-Way Quantum Computer*, Physical Review Letters **86**, 5188 (2001), 0510135v1.

[12] S. Morley-Short, S. Bartolucci, M. Gimeno-Segovia, P. Shadbolt, H. Cable, and T. Rudolph, *Physical-depth architectural requirements for generating universal photonic cluster states*, Quantum Science and Technology **3**, 015005 (2018), 1706.07325.

[13] R. Raussendorf, D. E. Browne, and H. J. Briegel, *Measurement-based quantum computation on cluster states*, Physical Review A **68**, 022312 (2003), 0301052.

[14] D. Gottesman, *The Heisenberg Representation of Quantum Computers*, **1**, 1 (1998), 9807006.

[15] S. J. Devitt, A. G. Fowler, A. M. Stephens, A. D. Greentree, L. C. L. Hollenberg, W. J. Munro, and K. Nemoto, *Architectural design for a topological cluster state quantum computer*, New Journal of Physics **11** (2009), 0808.1782.

[16] R. Babbush, P. J. Love, and A. Aspuru-Guzik, *Adiabatic Quantum Simulation of Quantum Chemistry*, Scientific Reports **4**, 6603 (2015).

[17] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *Quantum Computation by Adiabatic Evolution*, (2000), 0001106.

[18] A. Aspuru-Guzik, *Simulated Quantum Computation of Molecular Energies*, Science **309**, 1704 (2005).

[19] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem*, Science **292**, 472 (2001), 0104129.

[20] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, *Simulation of electronic structure Hamiltonians using quantum computers*, Molecular Physics **109**, 735 (2011), 1001.3855.

[21] G. C. Knee and W. J. Munro, *Optimal Trotterization in universal quantum simulators under faulty control*, Physical Review A **91**, 052327 (2015).

[22] A. O. Niskanen, K. Harrabi, F. Yoshihara, Y. Nakamura, S. Lloyd, and J. S. Tsai, *Quantum Coherent Tunable Coupling of Superconducting Qubits*, Science **316**, 723 (2007).

[23] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, *An elementary quantum network of single atoms in optical cavities*, Nature **484**, 195 (2012), 1202.5955.

[24] C. J. Ballance *et al.*, *Hybrid quantum logic and a test of Bell's inequality using two different atomic isotopes*, Nature **528**, 384 (2015), 1505.04014.

[25] M. V. G. Dutt, L. Childress, L. Jiang, E. Togan, J. Maze, F. Jelezko, A. S. Zibrov, P. R. Hemmer, and M. D. Lukin, *Quantum register based on individual electronic and nuclear spin qubits in diamond.*, Science (New York, N.Y.) **316**, 1312 (2007).

[26] D. Press, T. D. Ladd, B. Zhang, and Y. Yamamoto, *Complete quantum control of a single quantum dot spin using ultrafast optical pulses*, Nature **456**, 218 (2008).

[27] D. Gottesman, *A Theory of Fault-Tolerant Quantum Computation*, Physical Review A **57**, 30 (1997), 9702029.

[28] A. M. Steane, *Efficient fault-tolerant quantum computing*, Nature **399**, 21 (1998), 9809054.

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

[30] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger, *Blueprint for a microwave trapped ion quantum computer*, Science Advances **3**, e1601540 (2017), 1508.00420.

[31] M. Riebe *et al.*, *Deterministic quantum teleportation with atoms.*, Nature **429**, 734 (2004).

[32] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, *Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects*, Physical Review A **89**, 022317 (2014), 1208.0391.

[33] 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), 1709.06952.

[34] 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**, 1 (2016), 1512.04600.

[35] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, *Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links*, Physical Review X **4**, 1 (2014), 1406.0880.

[36] I. A. Walmsley and J. Nunn, *Editorial: Building Quantum Networks*, Physical Review Applied **6**, 040001 (2016).

[37] A. Gheorghiu, E. Kashefi, and P. Wallden, *Robustness and device independence of verifiable blind quantum computing*, New Journal of Physics **17**, 083040 (2015), 1502.02571.

[38] W. J. 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), 0401076.

[39] M. Giustina *et al.*, *Bell violation using entangled photons without the fair-sampling assumption*, Nature **497**, 227 (2013), 1309.0712v2.

[40] B. G. Christensen *et al.*, *Detection-loophole-free test of quantum nonlocality, and applications*, Physical Review Letters **111**, 1 (2013), 1306.5772.

[41] L. K. Shalm *et al.*, *Strong Loophole-Free Test of Local Realism*, Physical Review Letters **115**, 1 (2015), 1511.03189.

[42] F. Bussières, 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).

[43] X. L. Wang *et al.*, *Experimental ten-photon entanglement*, Physical Review Letters **117**, 1 (2016), 1605.08547.

[44] H. Bernien *et al.*, *Probing many-body dynamics on a 51-atom quantum simulator*, Nature **551**, 579 (2017), 1707.04344.

[45] J. Kelly, *A Preview of Bristlecone, Google's New Quantum Processor*, 2018, https://research.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html.

[46] W. K. Wootters and W. H. Zurek, *A single quantum cannot be cloned*, Nature **299**, 802 (1982), 1511.00407.

[47] D. Dieks, *Communication by EPR devices*, Physics Letters A **92**, 271 (1982).

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

[49] P. Kok, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, *Linear optical quantum computing with photonic qubits*, Reviews of Modern Physics **79**, 135 (2007).

[50] Y. Li, P. C. Humphreys, G. J. Mendoza, and S. C. Benjamin, *Resource costs for fault-tolerant linear optical quantum computing*, Physical Review X **5**, 1 (2015), 1504.02457.

[51] A. Javadi *et al.*, *Single-photon non-linear optics with a quantum dot in a waveguide*, Nature Communications **6**, 8655 (2015), arXiv:1504.06895v1.

[52] S. Sun, H. Kim, G. S. Solomon, and E. Waks, *A quantum phase switch between a single solid-state spin and a photon*, Nature Nanotechnology **11**, 539 (2016), 1506.06036.

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

[54] 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).

[55] 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), 1411.6758.

[56] S. J. Devitt, A. D. Greentree, A. M. Stephens, and R. Van Meter, *High-speed quantum networking by ship*, Scientific Reports **6**, 1 (2016), 1410.3224.

[57] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Theoretical Computer Science , 175 (2011).

[58] J. Yin *et al.*, *Satellite-to-Ground Entanglement-Based Quantum Key Distribution*, Physical Review Letters **119**, 200501 (2017).

[59] S.-K. Liao *et al.*, *Satellite-to-ground quantum key distribution*, Nature **549**, 43 (2017), 1707.00542.

[60] K. Boone, J.-P. Bourgoin, E. Meyer-Scott, K. Heshami, T. Jennewein, and C. Simon, *Entanglement over global distances via quantum repeaters with satellite links*, Physical Review A **91**, 052325 (2015), 1410.5384.

[61] 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), 9803056v1.

[62] M. Zukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert, *"Event-ready-detectors" Bell experiment via entanglement swapping*, Physical Review Letters **71**, 4287 (1993).

[63] S.-W. Lee and H. Jeong, *Bell-state measurement and quantum teleportation using linear optics: two-photon pairs, entangled coherent states, and hybrid entanglement*, (2013), 1304.1214.

[64] 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**, 1 (2007), 0701239.

[65] A. Kuhn, M. Hennrich, and G. Rempe, *Deterministic Single-Photon Source for Distributed Quantum Networking*, Physical Review Letters **89**, 4 (2002), 0204147.

[66] H. G. Barros, A. Stute, T. E. Northup, C. Russo, P. O. Schmidt, and R. Blatt, *Deterministic single-photon source from a single ion*, New Journal of Physics **11** (2009), 0905.2885.

[67] I. Schwartz, D. Cogan, E. R. Schmidgall, Y. Don, L. Gantz, O. Kenneth, N. H. Lindner, and D. Gershoni, *Deterministic generation of a cluster state of entangled photons*, Science **354**, 434 (2016), 1606.07492.

[68] P.-I. Schneider, N. Srocka, S. Rodt, L. Zschiedrich, S. Reitzenstein, and S. Burger, *Numerical optimization of the extraction efficiency of a quantum-dot based single-photon emitter into a single-mode fiber*, Optics Express **26**, 8479 (2018), 1803.03438.

[69] N. Somaschi *et al.*, *Near-optimal single-photon sources in the solid state*, Nature Photonics **10**, 340 (2016), 1510.06499.

[70] A. Christ and C. Silberhorn, *Limits on the deterministic creation of pure single-photon states using parametric down-conversion*, Physical Review A **85**, 023829 (2012), 1111.4095.

[71] A. Migdall, D. Branning, and S. Castelletto, *Tailoring single-photon and multiphoton probabilities of a single-photon on-demand source*, Physical Review A **66**, 053805 (2002).

[72] X.-S. Ma, S. Zotter, J. Kofler, T. Jennewein, and A. Zeilinger, *Experimental generation of single photons via active multiplexing*, Physical Review A **83**, 043814 (2011).

[73] C. Joshi, A. Farsi, S. Clemmen, S. Ramelow, and A. L. Gaeta, *Frequency multiplexing for quasi-deterministic heralded single-photon sources*, Nature Communications **9**, 847 (2018).

[74] J. Nunn, N. Langford, W. Kolthammer, T. F. M. Champion, M. Sprague, P. Michelberger, X.-M. Jin, D. England, and I. A. Walmsley, *Enhancing Multiphoton Rates with Quantum Memories*, Physical Review Letters **110**, 133601 (2013).

[75] G. J. Mendoza, R. Santagati, J. Munns, E. Hemsley, M. Piekarek, E. Martín-López, G. D. Marshall, D. Bonneau, M. G. Thompson, and J. L. O'Brien, *Active temporal and spatial multiplexing of photons*, Optica **3**, 127 (2016), 1503.01215.

[76] 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), 0906.2699.

[77] Z. Zheng, O. Mishina, N. Treps, and C. Fabre, *Atomic quantum memory for multimode frequency combs*, Physical Review A **91**, 031802 (2015).

[78] S. Massar and S. Popescu, *Optimal extraction of information from finite quantum ensembles*, Physical Review Letters **74**, 1259 (1995).

[79] M. Ho, J.-D. Bancal, and V. Scarani, *Device-independent certification of the teleportation of a qubit*, Physical Review A **88**, 052318 (2013), 1308.0084v1.

[80] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, *Proposed Experiment to Test Local Hidden-Variable Theories*, Physical Review Letters **23**, 880 (1969).

[81] P. Hyllus, O. Gühne, D. Bruß, and M. Lewenstein, *Relations between entanglement witnesses and Bell inequalities*, Physical Review A **72**, 012321 (2005), 0504079.

[82] P. W. Shor, *Scheme for reducing decoherence in quantum computer memory*, Physical Review A **52**, 2493 (1995), 0506097.

[83] 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), 1501.03980v2.

[84] T. Chanelière, G. Hétet, and N. Sangouard, *Quantum optical memory protocols in atomic ensembles*, (2018), 1801.10023.

[85] 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), 1704.00879.

[86] C. Langer *et al.*, *Long-Lived Qubit Memory Using Atomic Ions*, Physical Review Letters **95**, 060502 (2005).

[87] 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).

[88] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin, *Multimode quantum memory based on atomic frequency combs*, Physical Review A **79**, 052329 (2009), 0805.4164.

[89] P. C. Maurer *et al.*, *Room-temperature quantum bit memory exceeding one second.*, Science (New York, N.Y.) **336**, 1283 (2012).

[90] C. Simon *et al.*, *Quantum memories :A review based on the European integrated project"Qubit Applications (QAP)"*, European Physical Journal D **58**, 1 (2010), 1003.1107.

[91] E. Hahn, *Spin Echoes*, Physical Review **80**, 580 (1950).

[92] I. D. Abella, N. A. Kurnit, and S. R. Hartmann, *Photon Echoes*, Physical Review **141**, 391 (1966).

[93] 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).

[94] B. Lauritzen, J. Miná, H. de Riedmatten, M. Afzelius, N. Sangouard, C. Simon, and N. Gisin, *Telecommunication-Wavelength Solid-State Memory at the Single Photon Level*, Physical Review Letters **104**, 080502 (2010), 0908.2348.

[95] 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), 0612169.

[96] G. Hétet, M. Hosseini, B. M. Sparkes, D. Oblak, P. K. Lam, and B. C. Buchler, *Photon echoes generated by reversing magnetic field gradients in a rubidium vapor*, Optics Letters **33**, 2323 (2008), 0806.4258.

[97] G. Hétet, J. J. Longdell, M. J. Sellars, P. K. Lam, and B. C. Buchler, *Multimodal Properties and Dynamics of Gradient Echo Quantum Memory*, Physical Review Letters **101**, 203601 (2008), 0801.3860.

[98] Y.-W. Cho, G. T. Campbell, J. L. Everett, J. Bernu, D. B. Higginbottom, M. T. Cao, J. Geng, N. P. Robins, P. K. Lam, and B. C. Buchler, *Highly efficient optical quantum memory with long coherence time in cold atoms*, Optica **3**, 100 (2016), 1601.04267.

[99] M. Fleischhauer and M. D. Lukin, *Dark-State Polaritons in Electromagnetically Induced Transparency*, Physical Review Letters **84**, 5094 (2000), 0001094.

[100] 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), 0604037.

[101] I. Novikova, N. B. Phillips, and A. V. Gorshkov, *Optimal light storage with full pulse-shape control*, Physical Review A **78**, 021802 (2008), 0805.1927v1.

[102] 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).

[103] S. Manz, T. Fernholz, J. Schmiedmayer, and J.-W. Pan, *Collisional decoherence during writing and reading quantum states*, Physical Review A **75**, 040101 (2007).

[104] D. S. Ding, K. Wang, W. Zhang, S. Shi, M. X. Dong, Y. C. Yu, Z. Y. Zhou, B. S. Shi, and G. C. Guo, *Entanglement between low- and high-lying atomic spin waves*, Physical Review A **94**, 052326 (2016).

[105] 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).

[106] J. Nunn, I. Walmsley, M. Raymer, K. Surmacz, F. Waldermann, Z. Wang, and D. Jaksch, *Mapping broadband single-photon wave packets into an atomic memory*, Physical Review A **75**, 011401 (2007).

[107] 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).

[108] D. England, P. Bustard, J. Nunn, R. Lausten, and B. Sussman, *From Photons to Phonons and Back: A THz Optical Memory in Diamond*, Physical Review Letters **111**, 243601 (2013).

[109] 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).

[110] K. Surmacz, J. Nunn, K. Reim, K. C. Lee, V. O. Lorenz, B. Sussman, I. A. Walmsley, and D. Jaksch, *Efficient spatially resolved multimode quantum memory*, Physical Review A **78**, 033806 (2008).

[111] M. Parniak, M. Dąbrowski, M. Mazelanik, A. Leszczyński, M. Lipka, and W. Wasilewski, *Wavevector multiplexed atomic quantum memory via spatially-resolved single-photon detection*, Nature Communications **8**, 2140 (2017), 1706.04426.

[112] 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), 0105105.

[113] A. I. Lvovsky, B. C. Sanders, and W. Tittel, *Optical quantum memory*, Nature Photonics **3**, 706 (2009), arXiv:1002.4659v3.

[114] M. D. Eisaman, L. Childress, A. André, F. Massou, A. S. Zibrov, and M. D. Lukin, *Shaping Quantum Pulses of Light Via Coherent Atomic Memory*, Physical Review Letters **93**, 233602 (2004), 0406093.

[115] K. T. Kaczmarek *et al.*, *A high-speed noise-free optical quantum memory*, Physical Review A **97**, 042316 (2018), 1704.00013.

[116] R. Finkelstein, E. Poem, O. Michel, O. Lahad, and O. Firstenberg, *Fast, noise-free memory for photon synchronization at room temperature*, Science Advances **4**, eaap8598 (2018), 1708.01919.

[117] K. T. Kaczmarek, *Towards an integrated noise-free Quantum Memory*, PhD thesis, University of Oxford, 2017.

[118] P. S. Michelberger *et al.*, *Interfacing GHz-bandwidth heralded single photons with a warm vapour Raman memory*, New Journal of Physics **17**, 043006 (2015).

[119] 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), 1509.05098.

[120] B. C. Buchler, M. Hosseini, G. Hétet, B. M. Sparkes, and P. K. Lam, *Precision spectral manipulation of optical pulses using a coherent photon echo memory*, Optics Letters **35**, 1091 (2010), 1202.6096.

[121] M. Hosseini, B. M. Sparkes, G. Campbell, P. K. Lam, and B. C. Buchler, *High efficiency coherent optical memory with warm rubidium vapour.*, Nature communications **2**, 174 (2011), 1009.0567.

[122] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler, *Unconditional room-temperature quantum memory*, Nature Physics **7**, 794 (2011).

[123] M. P. Hedges, J. J. Longdell, Y. Li, and M. J. Sellars, *Efficient quantum memory for light*, Nature **465**, 1052 (2010).

[124] C. Clausen, I. Usmani, F. Bussiéres, N. Sangouard, M. Afzelius, H. De Riedmatten, and N. Gisin, *Quantum storage of photonic entanglement in a crystal*, Nature **469**, 508 (2011), 1009.0489.

[125] E. Saglamyurek, N. Sinclair, J. Jin, J. A. Slater, D. Oblak, W. Tittel, F. Bussières, M. George, R. Ricken, and W. Sohler, *Broadband waveguide quantum memory for entangled photons (Invited Paper)*, Nature Letters **469**, 512 (2011), 1009.0490.

[126] 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).

[127] P. Jobez, C. Laplane, N. Timoney, N. Gisin, A. Ferrier, P. Goldner, and M. Afzelius, *Coherent Spin Control at the Quantum Level in an Ensemble-Based Optical Memory*, Physical Review Letters **114**, 230502 (2015).

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

[129] 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), 0511014.

[130] 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).

[131] 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).

[132] M. Hennrich, T. Legero, A. Kuhn, and G. Rempe, *Vacuum-stimulated raman scattering based on adiabatic passage in a high-finesse optical cavity*, Physical review letters **85**, 4872 (2000).

[133] J. Nunn, *Quantum Memory in Atomic Ensembles*, PhD thesis, University of Oxford, 2008.

[134] 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).

[135] M. R. Sprague, P. S. Michelberger, T. F. M. Champion, D. G. England, J. Nunn, X.-M. Jin, W. S. Kolthammer, A. Abdolvand, P. S. J. Russell, and I. A. Walmsley, *Broadband single-photon-level memory in a hollow-core photonic crystal fibre*, Nature Photonics **8**, 287 (2014).

[136] M. Dąbrowski, R. Chrapkiewicz, and W. Wasilewski, *Hamiltonian design in readout from room-temperature Raman atomic memory*, Optics Express **22**, 26076 (2014), 1406.6489.

[137] R. A. Campos, B. E. A. Saleh, and M. C. Teich, *Quantum-Mechanical Lossless Beam Splitter: SU(2) Symmetry and Photon Statistics*, Physical Review A **40** (1989).

[138] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. I. Cavity model*, Physical Review A **76**, 033804 (2007).

[139] 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).

[140] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. III. Effects of inhomogeneous broadening*, Physical Review A **76**, 033806 (2007), 0612084.

[141] A. V. Gorshkov, T. Calarco, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. IV. Optimal control using gradient ascent*, Physical Review A **77**, 043806 (2008).

[142] 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), 1610.03743.

[143] D. A. Steck, *Cesium D Line Data*, 2003, http://steck.us/alkalidata.

[144] R. Loudon, *The Quantum Theory of Light*, Third edition (OUP Oxford, 2000).

[145] I. Vurgaftman and M. Bashkansky, *Suppressing four-wave mixing in warm-atomic-vapor quantum memory*, Physical Review A **87**, 063836 (2013).

[146] 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.

[147] J. Nunn *et al.*, *Theory of noise suppression in Λ-type quantum memories by means of a cavity*, Physical Review A **96**, 012338 (2016), 1601.00157.

[148] 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*, Phys. Rev. B **91**, 205108 (2015).

[149] H. Wu and M. Xiao, *Cavity linewidth narrowing and broadening due to competing linear and nonlinear dispersions*, Optics Letters **32**, 3122 (2007).

[150] M. G. Raymer and C. J. McKinstrie, *Quantum input-output theory for optical cavities with arbitrary coupling strength: Application to two-photon wave-packet shaping*, Physical Review A **88**, 043819 (2013).

[151] 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 multi-photon statistics*, Physical Review A **88**, 013822 (2013).

[152] A. Christ, K. Laiho, A. Eckstein, K. N. Cassemiro, and C. Silberhorn, *Probing multimode squeezing with correlation functions*, New Journal of Physics **13** (2011), 1012.0262.

[153] 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, *A Cavity-Enhanced Room-Temperature Broadband Raman Memory*, Physical Review Letters **116**, 090501 (2015), 1510.04625.

[154] J. H. D. Munns, C. Qiu, P. M. Ledingham, I. A. Walmsley, J. Nunn, and D. J. Saunders, *In situ characterization of an optically thick atom-filled cavity*, Physical Review A **93**, 013858 (2016), 1511.05448.

[155] T. Hansch and B. Couillaud, *Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity*, Optics Communications **35**, 441 (1980).

[156] B. Beauregard, *Improving the Beginner's PID*, 2011, http://brettbeauregard.com/blog/2011/04/improving-the-beginners-pid-introduction/.

[157] P. Siddons, C. S. Adams, and I. G. Hughes, *Off-resonance absorption and dispersion in vapours of hot alkali-metal atoms*, Journal of Physics B: Atomic, Molecular and Optical Physics **42**, 175004 (2009).

[158] M. A. Zentile, J. Keaveney, L. Weller, D. J. Whiting, C. S. Adams, and I. G. Hughes, *ElecSus: A program to calculate the electric susceptibility of an atomic ensemble*, Computer Physics Communications **189**, 162 (2015).

[159] P. Siddons, C. S. Adams, C. Ge, and I. G. Hughes, *Absolute absorption on rubidium D lines: comparison between theory and experiment*, Journal of Physics B: Atomic, Molecular and Optical Physics **41**, 155004 (2008).

[160] G. A. Pitz, C. D. Fox, and G. P. Perram, *Pressure broadening and shift of the cesium D_2 transition by the noble gases and N_2, H_2, HD, D_2, CH_4, C_2H_6, CF_4, and ^3He with comparison to the D_1 transition*, Physical Review A **82**, 042502 (2010).

[161] O. Kozlova, S. Guérandel, and E. de Clercq, *Temperature and pressure shift of the Cs clock transition in the presence of buffer gases: Ne, N_2, Ar*, Physical Review A **83**, 062714 (2011).

[162] P. Thompson, D. E. Cox, and J. B. Hastings, *Rietveld Refinement of Debye-Scherrer Synchrotron X-ray Data from Al_2O_3*, J. Appl. Cryst. **20**, 79 (1987).

[163] R. J. Wells, *Rapid approximation to the Voigt/Faddeeva function and its derivatives*, Journal of Quantitative Spectroscopy and Radiative Transfer **62**, 29 (1999).

[164] S. G. Johnson, *Faddeeva Package*, 2012, http://ab-initio.mit.edu/wiki/index.php/Faddeeva_Package.

[165] A. Wicht, M. Muller, R.-H. Rinkleff, A. Rocco, and K. Danzmann, *Experimental demonstration of negative dispersion without absorption*, Optics Communications **179**, 107 (2000).

[166] K. F. Reim, *Broadband optical quantum memory*, PhD thesis, University of Oxford, 2011.

[167] L. Weller, R. J. Bettles, P. Siddons, C. S. Adams, and I. G. Hughes, *Absolute absorption on the rubidium D 1 line including resonant dipoledipole interactions*, Journal of Physics B: Atomic, Molecular and Optical Physics **44**, 195006 (2011).

[168] 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).

[169] M. R. Sprague, *Raman Memory for Entanglement in Diamonds and Light Storage in Optical Fibres*, PhD thesis, University of Oxford, 2014.

[170] K. F. Reim, J. Nunn, X.-M. Jin, P. S. Michelberger, T. F. M. Champion, D. G. England, K. C. Lee, W. S. Kolthammer, N. K. Langford, and I. a. Walmsley, *Multipulse Addressing of a Raman Quantum Memory: Configurable Beam Splitting and Efficient Readout*, Physical Review Letters **108**, 263602 (2012).

[171] A. B. Matsko, I. Novikova, M. O. Scully, and G. R. Welch, *Radiation trapping in coherent media*, Physical Review Letters **87**, 1 (2001), 0101147.

[172] M. Parniak and W. Wasilewski, *Direct observation of atomic diffusion in warm rubidium ensembles*, Applied Physics B: Lasers and Optics **116**, 415 (2014), 1308.2528.

[173] W. Happer, *Optical Pumping*, Reviews of Modern Physics **44**, 169 (1972).

[174] H. Kogelnik and T. Li, *Laser beams and resonators*, Proceedings of the IEEE **54**, 1312 (1966), arXiv:1003.1107v1.

[175] 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), 0703111.

[176] C. Wu, M. G. Raymer, Y. Y. Wang, and F. Benabid, *Quantum theory of phase correlations in optical frequency combs generated by stimulated Raman scattering*, Physical Review A **82**, 053834 (2010).

[177] S. Groeger, G. Bison, J. L. Schenker, R. Wynands, and A. Weis, *A high-sensitivity laser-pumped Mx magnetometer*, European Physical Journal D **38**, 239 (2006), 0406105v2.

[178] A. Horsley, *High Resolution Field Imaging with Atomic Vapor Cells*, PhD thesis, University of Basel, 2015.

[179] E. P. Corsini, T. Karaulanov, M. Balabas, and D. Budker, *Hyperfine frequency shift and Zeeman relaxation in alkali-metal-vapor cells with antirelaxation alkene coating*, Physical Review A - Atomic, Molecular, and Optical Physics **87**, 1 (2013).

[180] P. K. Lam, *Multiparameter Optimisation of a Magneto-Optical Trap Using Deep Learning Artificial Neural Networks*, International Conference on Quantum Communication Measurement and Computing (2018).

[181] M. Parniak, D. Pecak, and W. Wasilewski, *Multimode Raman light-atom interface in warm atomic ensemble as multiple three-mode quantum operations*, (2015), 1505.041118.

[182] G. Adenier, D. Calonico, S. Micalizio, N. Samantaray, I. P. Degiovanni, and I. R. Berchera, *Realization of a twin beam source based on four-wave mixing in Cesium*, (2016), 1604.01807.

[183] 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), 1706.03801.

[184] 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), 1504.06251v1.

[185] V. Ansari, J. M. Donohue, B. Brecht, and C. Silberhorn, *Tailoring nonlinear processes for quantum optics with pulsed temporal-mode encodings*, (2018), 1803.04316.

[186] S. Xu, H. Z. Shen, and X. X. Yi, *Demultiplexing of photonic temporal modes by a linear system*, Physical Review A **97**, 033841 (2018).

[187] M. Mirhosseini, O. S. Magaña-Loaiza, M. N. O'Sullivan, B. Rodenburg, M. Malik, M. P. Lavery, M. J. Padgett, D. J. Gauthier, and R. W. Boyd, *High-dimensional quantum cryptography with twisted light*, New Journal of Physics **17**, 1 (2015), 1402.7113.

[188] M. Krenn, J. Handsteiner, M. Fink, R. Fickler, and A. Zeilinger, *Twisted photon entanglement through turbulent air across Vienna*, Proceedings of the National Academy of Sciences **112**, 14197 (2015), 1507.06551.

[189] P. C. Humphreys, W. S. Kolthammer, J. Nunn, M. Barbieri, A. Datta, and I. A. Walmsley, *Continuous-Variable Quantum Computing in Optical Time-Frequency Modes Using Quantum Memories*, Physical Review Letters **113**, 130502 (2014).

[190] V. A. Averchenko, V. Thiel, and N. Treps, *Nonlinear photon subtraction from a multimode quantum field*, Physical Review A **89**, 063808 (2014), 1404.0160.

[191] 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).

[192] V. Wong and I. A. Walmsley, *Linear filter analysis of methods for ultrashort-pulse-shape measurements*, Journal of the Optical Society of America B **12**, 1491 (1995).

[193] 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), 1002.0350.

[194] A. Eckstein, B. Brecht, and C. Silberhorn, *A quantum pulse gate based on spectrally engineered sum frequency generation*, Optics Express **19**, 13770 (2011), 1007.3215.

[195] B. Brecht, A. Eckstein, A. Christ, H. Suche, and C. Silberhorn, *From quantum pulse gate to quantum pulse shaperengineered frequency conversion in nonlinear optical waveguides*, New Journal of Physics **13**, 065029 (2011), 1101.6060v1.

[196] D. V. Reddy, M. G. Raymer, C. J. McKinstrie, L. Mejling, and K. Rottwitt, *Temporal mode selectivity by frequency conversion in second-order nonlinear optical waveguides*, Optics Express **21**, 13840 (2013), 1305.2165.

[197] 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), 1610.08326.

[198] K. Tordrup, A. Negretti, and K. Mølmer, *Holographic quantum computing*, Physical Review Letters **101**, 1 (2008), 0802.4406.

[199] D. V. Vasilyev, I. V. Sokolov, and E. S. Polzik, *Quantum memory for images: A quantum hologram*, Physical Review A **77**, 020302 (2008), 0704.1737.

[200] J. Schwinger, *Unitary Operator Bases.*, Proceedings of the National Academy of Sciences of the United States of America **46**, 570 (1960).

[201] M. Mohseni, A. T. Rezakhani, and D. A. Lidar, *Quantum-process tomography: Resource analysis of different strategies*, Physical Review A **77**, 032322 (2008), 0702131.

[202] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd edition (Cambridge University Press, 2007).

[203] K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering*, 3rd edition (Cambridge University Press, 2006).

[204] S. Brierley, S. Weigert, and I. Bengtsson, *All Mutually Unbiased Bases in Dimensions Two to Five*, (2009), 0907.4097.

[205] A. Monmayrant, S. Weber, and B. Chatel, *A newcomer's guide to ultrashort pulse shaping and characterization*, Journal of Physics B: Atomic, Molecular and Optical Physics **43**, 103001 (2010).

[206] Z. Hradil, J. Řeháček, J. Fiurášek, and M. Ježek, *Maximum-Likelihood Methods in Quantum Mechanics*, Lect. Notes Phys. **649**, 59 (2004).

[207] J. Řeháček, Z. Hradil, E. Knill, and A. I. Lvovsky, *Diluted maximum-likelihood algorithm for quantum tomography*, Physical Review A **75**, 042108 (2007), 0611244.

[208] Z. Hradil, J. Summhammer, and H. Rauch, *Quantum tomography as normalization of incompatible observations*, Physics Letters A **261**, 20 (1999), 9806014.

[209] Y. S. Teo, J. Řeháček, and Z. Hradil, *Informationally incomplete quantum tomography*, Quantum Measurements and Quantum Metrology **1**, 23 (2013), 1309.2906.

[210] R. Grobe, K. Rzazewski, and J. H. Eberly, *Measure of electron-electron correlation in atomic physics*, Journal of Physics B: Atomic, Molecular and Optical Physics **27** (1994).

[211] 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), 1702.05687.

[212] 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), 1210.8342.

[213] N. Quesada and J. E. Sipe, *High efficiency in mode-selective frequency conversion*, Optics Letters **41**, 364 (2016), 1508.03361.

[214] F. Krumm, J. Sperling, and W. Vogel, *Multitime correlation functions in nonclassical stochastic processes*, Physical Review A **93**, 1 (2016), 1603.09555.

[215] N. Quesada and J. E. Sipe, *Effects of time ordering in quantum nonlinear optics*, Physical Review A **90**, 063840 (2014), arXiv:1407.6976v1.

[216] F. Bouchard, K. Heshami, D. England, R. Fickler, R. W. Boyd, B.-G. Englert, L. L. Sánchez-Soto, and E. Karimi, *Experimental investigation of quantum key distribution protocols with twisted photons*, p. 1 (2018), 1802.05773.

[217] P. M. Leung, W. J. Munro, K. Nemoto, and T. C. Ralph, *Spectral Effects of Strong Chi-2 Non-Linearity for Quantum Processing*, Physical Review A **79**, 042307 (2008), 0811.2038.

[218] 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), 1501.03216.

[219] 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), 1702.03336.

[220] C. K. Hong, Z. Y. Ou, and L. Mandel, *Measurement of subpicosecond time intervals between two photons by interference*, Physical Review Letters **59**, 2044 (1987).

[221] P. J. Mosley, J. S. Lundeen, B. J. Smith, P. Wasylczyk, A. B. U'Ren, C. Silberhorn, and I. A. Walmsley, *Heralded Generation of Ultrafast Single Photons in Pure Quantum States*, Physical Review Letters **100**, 133601 (2008), 0711.1054.

[222] F. Graffitti, D. Kundys, D. T. Reid, A. M. Brańczyk, and A. Fedrizzi, *Pure down-conversion photons through sub-coherence-length domain engineering*, Quantum Science and Technology **2**, 035001 (2017).

[223] E. Meyer-Scott, N. Montaut, J. Tiedau, L. Sansoni, H. Herrmann, T. J. Bartley, and C. Silberhorn, *Limits on the heralding efficiencies and spectral purities of spectrally filtered single photons from photon-pair sources*, Physical Review A **95**, 061803 (2017), 1702.05501.

[224] J. H. D. Munns, *Towards Noise Suppression in a Broadband Quantum Memory in Warm Caesium Vapour*, M.res. thesis, Imperial College London, 2014.

[225] J. S. Steinhart and S. R. Hart, *Calibration curves for thermistors*, Deep-Sea Research **15**, 497 (1968).

[226] D. Kane and R. Trebino, *Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating*, IEEE Journal of Quantum Electronics **29**, 571 (1993).

[227] C. Iaconis and I. A. Walmsley, *Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses.*, Optics letters **23**, 792 (1998).

[228] C. Bourassin-Bouchet, M. M. Mang, I. Gianani, and I. A. Walmsley, *Mutual interferometric characterization of a pair of independent electric fields.*, Optics letters **38**, 5299 (2013).

[229] S. Park, *Complex Temporal Shaping and Characterization of Laser Pulses for Quantum Memories*, M.phys., University of Oxford, 2017.

[230] F. Gao and L. Han, *Implementing the Nelder-Mead simplex algorithm with adaptive parameters*, Computational Optimization and Applications **51**, 259 (2012).

[231] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, *Convergence Properties of the Nelder–Mead Simplex Method in Low Dimensions*, SIAM Journal on Optimization **9**, 112 (1998).

[232] J. B. Altepeter, D. F. V. James, and P. G. Kwiat, *Quantum State Estimation*, Lect. Notes Phys. **649**, 113 (2004), 1403.2781v1.

[233] V. Bužek, *Quantum Tomography from Incomplete Data via MaxEnt Principle*, Lect. Notes Phys. **649**, 189 (2004).

[234] R. Blume-Kohout, *Hedged Maximum Likelihood Quantum State Estimation*, Physical Review Letters **105**, 200504 (2010), 1001.2029.

[235] C. A. Fuchs and R. Schack, *Unknown Quantum States and Operations, a Bayesian View*, Lect. Notes Phys. **649**, 147 (2004), 0404156v1.

[236] L. N. Trefethen, *Spectral Methods in MATLAB*, 1st edition (SIAM, Philadelphia, 2000).

APPENDIX A

EXPERIMENTAL CONTROL INFRASTRUCTURE

In this appendix, I discuss the elements of the control systems which I have developed for the experiments contained in this thesis: Firstly: A.1 cavity memory experiment, then in A.2 some of the control elements for the memory optimisation experiments, and lastly in A.3 I introduce the systems for the pulse carving experiments.

A.1 Cavity Memory Experiment: Control Systems

A.1.1 Temperature Stabilisation

As mentioned in the main text, stabilisation of the temperature of the vapour cell containing the caesium is critical in order to have control of the number of caesium atoms we interact with. Note that the key criteria here are that the temperature (i) remains stable, and (ii) we can vary the temperature over some range – typically between $50 - 100°C$. That is to say that the *absolute* value of the temperature is less important[1] and we can measure the atom number via spectroscopic means.

The vapour cell is wrapped with wire (Lakeshore QT–32) around the circumference of the cylinder, which was heated by driving a variable DC current through the wire. Four microthermistors (ATC Semitec) were mounted between the wire and the cell surface, and their resistance monitored to feedback on the current.

In the vapour cell, caesium is deposited at the coldest part of the cell, so a "cold finger" driven by a peltier element was mounted half way along the length of the cell in order to prevent caesium accumulating on the cell windows[2] Again the temperature here is not important provided that it is constant (driven at constant voltage), and is cooler that the windows (verified by visual inspection).

The resistance of the thermistors provides an estimate of the cell temperature, and is the signal which we aim to stabilise.

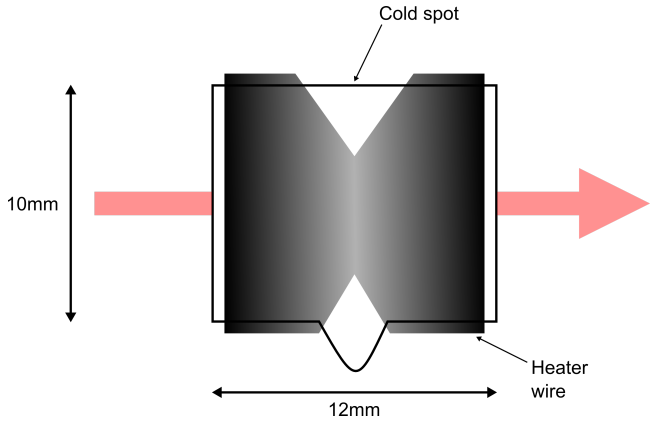
A.1.1(a) Thermistor Calibration

The four thermistors[3] were individually calibrated against Omega Precision fine wire thermocouples in order to compensate for any variation between them once wired up, such as due to a difference in

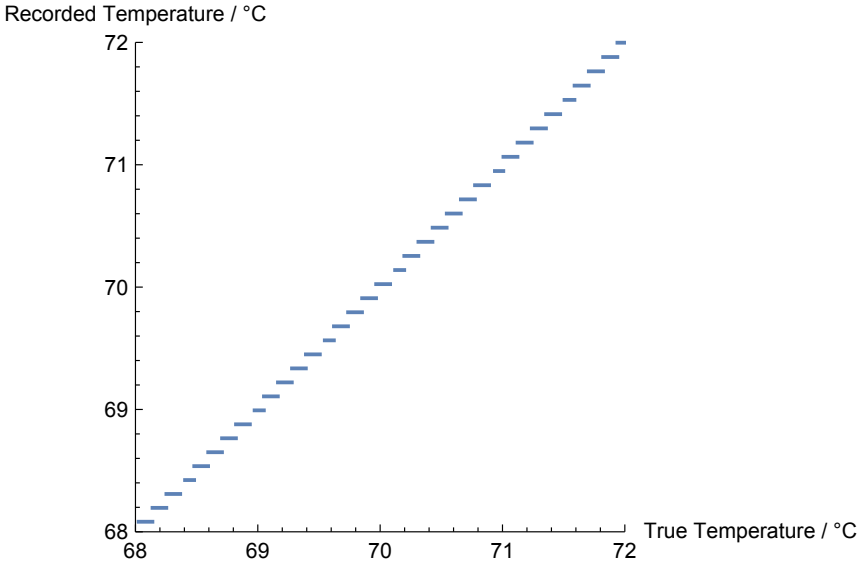
[1] Particularly since, as noted in section 3.3.4 that a single temperature measure is not sufficient to describe the operating conditions.

[2] Which could otherwise scatter light, and increase the intracavity loss. Cooling fins on the other face of the peltier we mounted outside the top of the box around the optical table to minimise thermal air currents disturbing the optics.

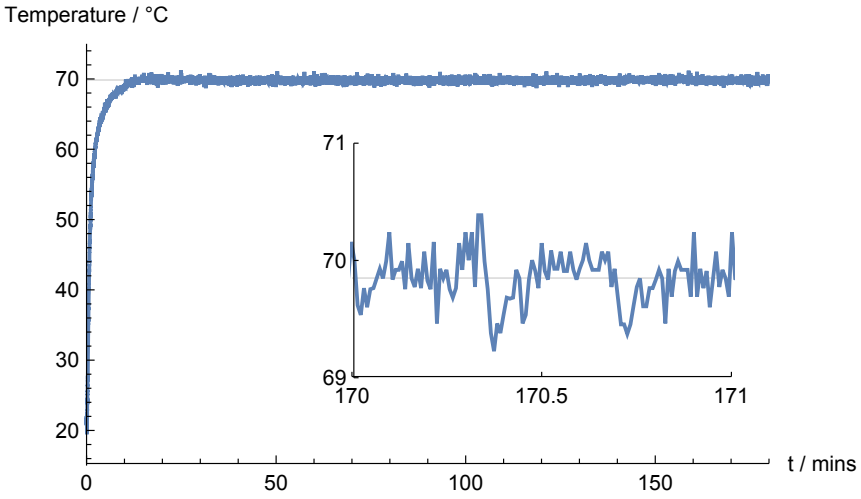
[3] For redundancy.



(a) Schematic of the vapour cell heating/cooling



(b) Temperature resolution as monitored by the Arduino MEGA 2560



(c) Thermal stability of the vapour cell as controlled with the Omega CN243 and monitored on the Arduino MEGA

Figure A.1: First implementation of cell temperature control and monitoring, figures from [224].

the length of their leads.

For calibration, each thermistor was taped directly to a thermocouple, and all of them fixed to a heater wrapped in aluminium foil. For each data point, the heater power supply was set to a specific value, and left to thermalize until fluctuations in the thermocouple reading were stable to within ±1K The data points for each thermistor was then fitted to a Steinhart-Hart [225] curve, defined by:

$$\frac{1}{T} = A + B\ln(R) + C\left[\ln(R)\right]^3 \tag{A.1}$$

These were wired to a circuit board in a potential divider configuration with a reference resistor – approximately matching the thermistor resistance at $\sim 70°\mathrm{C}$ to maximise sensitivity, and the voltage output monitored.

A.1.1(b) Heater Wire

The quad twisted heater wire was used to prevent stray magnetic fields from passing current, and increases the length of wire in contact with the cell by a factor of four.

A.1.1(c) Control Logic

In the first implementation, one of the thermistors was connected directly to an off-the-shelf temperature controller (Omega CN243) while the resistance of the remaining three were logged using an Arduino MEGA 2560 and written to file. The temperature controller implemented a simple on/off switch on the heater based upon the instantaneous resistance, where the heating current was provided by a variac.

The output of this system is shown in figures A.1 in which (b) shows the discretised temperature resulting from the Arduino MEGA 10-bit resolution, and (c) shows the logged temperature showing an on/off duty cycle on the order of 30s, fluctuating on the order of ±0.5°C.

To improve upon this, the microcontroller was upgraded to the DUE with 12-bit resolution (4096 levels) and with a faster clock speed, and all four thermistors monitored. These were then passed over serial to a LabView program in real time which implemented a simple PID algorithm, which feeds back to a variable power supply (Kiethley) such that the output voltage could be modulated. Once the PID settings were optimised, the improved resolution and continuous feedback enabled a stability of ±0.1°C.

A.1.2 Mechanical Stabilisation

A.1.2(a) Hänsch–Cauillaud locking

Our cavity has birefringence – necessarily for noise suppression scheme, so H and V polarisations acquire different roundtrip phase, and reflected light will be eliptically polarised.

We generate a locking signal by sending in a He:Ne laser at 633 nm which is far from the atomic resonance – and effectively doesn't see the atoms. This is prepared in diagonal polarisation (D) with respect to the cavity axes. Projecting the reflected light onto H and V using a PBS and monitoring the difference signal between the two ports we acquire information about the relative phase – and hence the cavity length.

The birefringent element, is a Lithium Niobate ($LiNbO_3$) crystal. This is uniaxial contributing a relative path length between fast and slow axes as a function of angle of incidence θ where the optic

(or "extraordinary") axis, e, is parallel to the axis of rotation, and the crystal rotates in the plane of the o axes.

$$\delta l(T,\theta) = l(T,\theta)\,[n_e(T) - n_o(T)] \tag{A.2}$$

where

$$l(T,\theta) = \frac{l(T_0)}{\cos(\theta)} \exp\left[\alpha(T - T_0)\right] \tag{A.3}$$

using the thermal expansion coefficient α, and

$$n_i(T) = n_i(T_0) + \frac{\mathrm{d}n_i}{\mathrm{d}T}\,(T - T_0) \quad \text{for } i \in \{o, e\} \tag{A.4}$$

With a combination of temperature and angle tuning we can achieve the required phase difference to optimise the cavity response for the signal and control polarisations.

The concave mirror of the cavity is mounted on a piezoelectric ring actuator which enabled feedback on the cavity length.

A.1.2(b) Locking Logic

The voltage of the two analysing photodiodes were monitored by a second Arduino (again originally MEGA2560, then upgraded to DUE for enhanced resolution & clock speed), this implemented a PID library D.J.S. and I adapted from [156].

Due to the difference in wavelength of the locking laser (633 nm) with the signal/control fields (852 nm), the periodicity of the locking signal is different to the cavity response at the signal frequency. To improve the longevity of the lock, we must remain within a single oscillation of the set-point. We therefore include a slow offset in the Arduino script which increments/decrements as the error signal approaches the edges of the signal window so as to follow slow drifts in the lab – such as the room temperature.

The output of the Arduino DAC channels have a voltage range spanning 0.55 V to 2.75 V, with 12–bit resolution (i.e. 4096 levels) the two outputs from the Arduino must be added in a ratio, then offset in voltage and finally amplified into the corresponding range for the piezo driver, such that the output range of the locking electronics spans the full range of the piezo driver. We choose a ratio 1:5,

In practice, the length is never explicitly evaluated by the lock. We scan the cavity length a range corresponding to a few cavity FSRs, and optimise on the fringe visibility by adjust the projection ratio onto the two photodiodes, and walking the HeNe alignment[4].

Once the visibility has been optimised, the piezo–driver is then switched to the locking set–up. The offset of the piezo can then be tuned via a potentiometer to find the initial set–point for the lock, for instance maximising the signal transmission. The lock was then engaged by a single–pole–double–throw switch, with the switch direction determining the sign of the feedback direction.[5] Once locked, the setpoint could be adjusted on the fly with a second potentiometer to maximise the transmission of the signal field through the cavity.

Whilst locked, the Arduino automatically increments/decrements the offset value if the error signal strays outside a tolerance region - thus enabling the cavity to remain locked to the same set–point for longer durations, typically on the order of an hour. The main loop of the Arduino locking is displayed

[4] The best visibility was obtained if the HeNe was misaligned with the cavity axis, and the second pass *only partially* overlaps with the reflected mode.

[5] Thus only locked successfully with 50% probability, each trial – so locking the cavity could take several attempts!

```
/*********************************************************************
*********************************************************************
  MAIN FUNCTION FOLLOWS:
*********************************************************************/
void loop() {

  outputRefresh(&Input, &outPIDDAC, true);

  double FracOutPID = OutPIDval/PIDmax;
  if ( abs(FracOutPID) > 0.8 ) {// increment/decrement offset:
    outOFFSETDAC = constrain(shiftOffset(outOFFSETDAC, OutPIDval), DACmin, DACmax);
  }

  if ( abs(FracOutPID) > 0.9 && myPID.GetMode() == AUTOMATIC ) {// wrap around as
      failsafe:
    myPID.SetMode(MANUAL);
    OutPIDval = -0.8*OutPIDval;
    myPID.SetMode(AUTOMATIC);
  }

  int SPDTstatTMP = SPDTread(inSPDT0, inSPDT1);
  if ( SPDTstatTMP != SPDTstat || SPDTstat == 0 ){// check offset/dir'n controls:
    SPDTstat = SPDTstatTMP;
    queryStatus(&Input, &Setpoint, &OutPIDval, &outPIDDAC, &outOFFSETDAC, &
        ControllerDirection, SPDTstat);
    myPID.SetControllerDirection(ControllerDirection);
  }

  int spPotDiff = analogRead(inPotSP) - spPotStat;
  if ( abs(spPotDiff) > potTolerance ){// check setpoint controls:
    spPotStat += spPotDiff;
    Setpoint = sp_update(Setpoint, spPotStat, sp0);
  }

}
```

Figure A.2: Main loop of Arduino locking code

in figure A.2.

We estimate the stability of the cavity, by monitoring the transmission of a probe laser through the cavity, locked to resonance, over an extended period of time. By post-selecting upon times when the laser frequency was within a narrow tolerance range ±5 MHz, and monitoring the variance of the *normalised* transmission (intensity referenced with respect to a pick off prior to the cavity) found to be on the order 1% as displayed in figure A.3, we determine the equivalent frequency uncertaintly in the lock via error propagation using the known cavity transmission function. This yields a frequency stability ~ 50 MHz, for our chosen PID settings.

The locking system parameters were continually monitored over serial, and graphically displayed on the master computer for the experiment, providing a warning if the lock had been lost.

A.1.3 Pump Laser Frequency Locking

I built a Doppler free saturation absorption spectroscopy set up to interface with the optical pumping ECDL controller (Toptica DLC), as shown in figure A.4 which enabled the diode laser (Toptica DL Pro) frequency to be locked to the atomic resonance which was used to reference the frequency scale of the scanning Fabry-Perot (Thorlabs) to determine the frequency offset of the Ti:Sa.

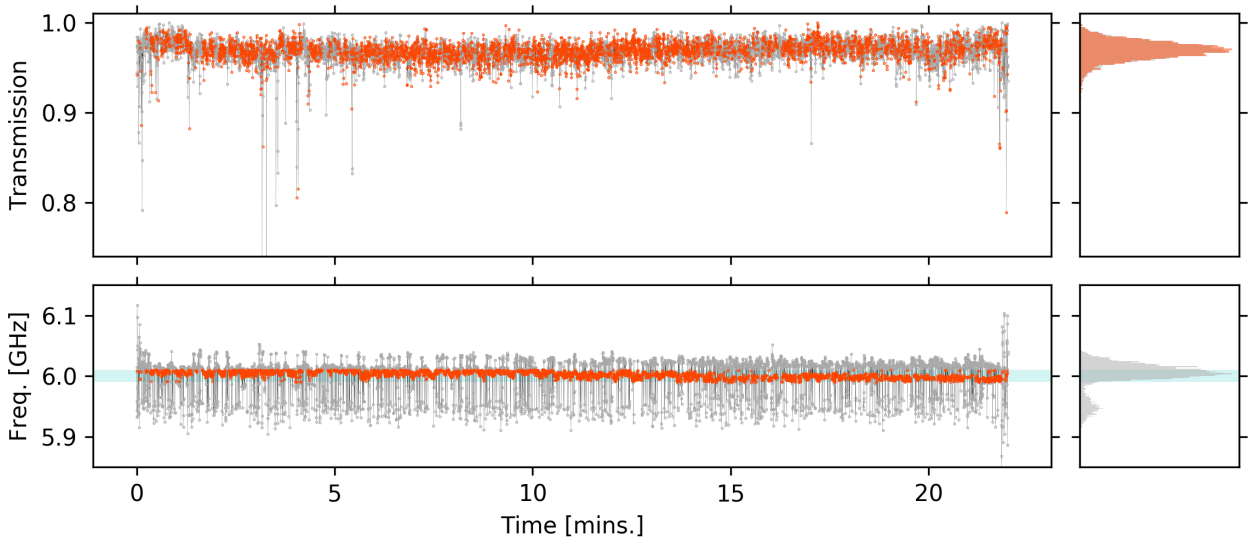


Figure A.3: Measurement of transmission through the cavity whilst locked. The orange points are those post-selected by when the Ti:Sa is within the frequency tolerance window – shaded blue. Note that here the Ti:Sa frequency is near a mode-jump and oscillates between two frequencies. Measurement performed by C.Q.

The Doppler free spectrum also provides a frequency reference for measurement of the optical depth and pumping efficiency of the atoms by scanning the laser across the transition frequencies, as in figure A.4.

A.2 Memory Optimisation Experiments

I wrote software in Python, to facilitate real-time automatic acquisition of memory data: this consisted of interfacing with three pololu servos to shutter off the three beams (signal, control, pump) in sequence of different configurations, and then acquire data from the oscilloscope which recorded the photodiode output after the memory. These were used for the acquisition of memory data in [142].

The oscilloscope libraries were also used for the automation of TEA+PARTI B and the downhill simplex search C.

A.3 Pulse Carving

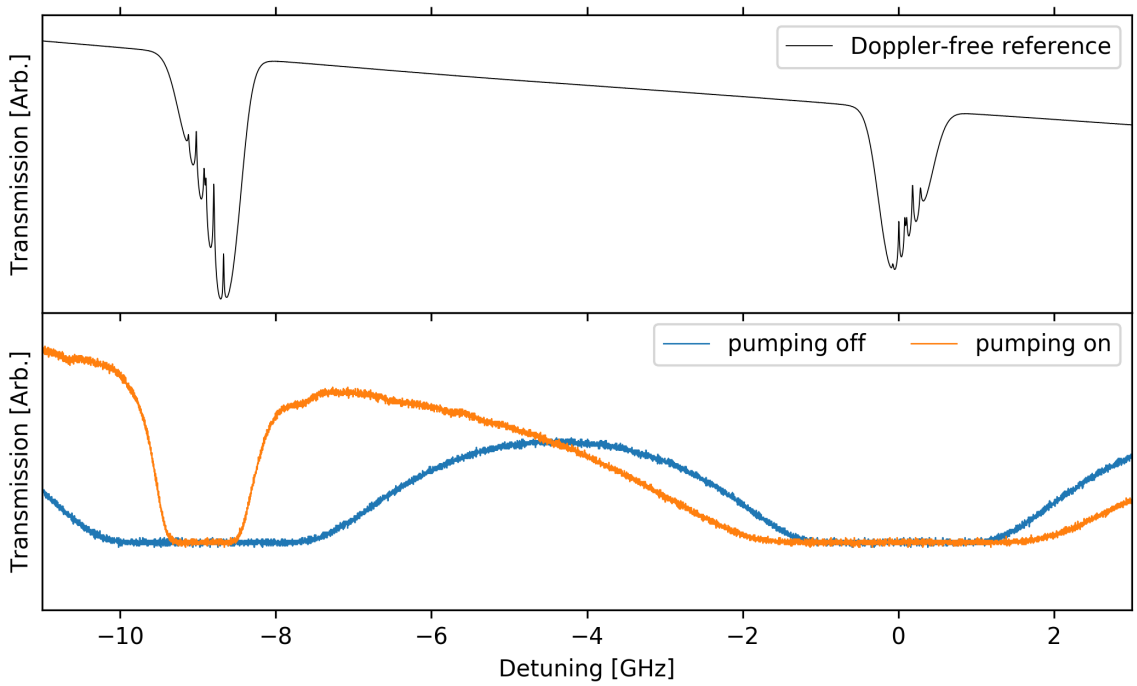
A.3.1 AWG Driver Software

I developed the control infrastructure for configuring, synchronising and running AWGs and pulse carving system.

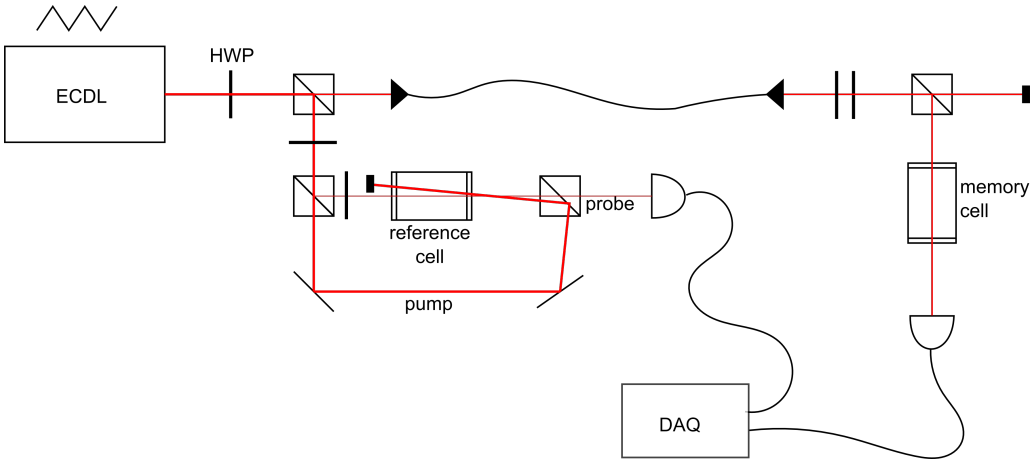
This consisted of Python libraries on the lab master computer communicating with the AWGs over the LAN, which ensured appropriate settings (e.g. clock speed, trigger settings). Additionally I wrote an interface for the evaluation of the necessary electrode waveforms starting from a desired electric field temporal profile.

This also facilitated the automation of electronically delaying the shape signal with respect with respect to the Ti:Sa, from which a reference clock (at 10 MHz) and trigger (TTL) was derived.

These libraries in addition to those I wrote for the oscilloscope enabled the automation of the downhill simplex search (appendix C) and TEA+PARTI (appendix B) measurements. This meant



(a) (top) Doppler free spectrum of D_2 transitions – see figure 3.4 for more detail, (bottom) Doppler broadened absorption spectrum, with optical pumping on/off. The intensity ramp when scanning frequency is present in both data (for these we extract $T \approx 78°C$, $w = -0.97$, pumping into $F = 3$ for measurements in part II of this thesis). Measurement performed by S.E.T.



(b) Set up for measurement of the bulk optical depth referenced by Doppler-free saturation absorption spectroscopy. The ECDL is driven by a sawtooth signal from [224].

Figure A.4: Bulk optical depth calibration and measurement. Single pass transmission of Doppler broadened profile can be used to extract vapour temperature and pumping efficiency.

that the output voltage/intensity could be recorded in real-time.

To illustrate the functionality of the control libraries – for the downhill search, this involved acquiring the output of a number of different waveforms, then the traces recorded from the 'scope were then used to update the waveform loaded to the AWGs for the next iteration (which would have been impractical by hand). This automated proceedure required that the AWGs and 'scope were synchronised appropriately such that the correct waveform had been written and had started (and was still) running before the oscilloscope acquired data to feed back into the algorithm. Since this algorithm had to be left running for several hours, the control infrastructure was tested to be robust over long periods, appropriately detecting and handling errors during the course of an experiment.

A.3.1(a) Electrode voltage configuration

As discussed in the main text (section 7.3.1) the response of the two electrodes was observed to be distinct, and we see that the finer features of the target waveform are not fully captured – seen in figure 7.1.

We adopted a manual optimisation proceedure to individually calibrate the settings of the two amplifiers. To achieve this we have control over the gain and crossing point. To configure both of these, two sinosoidal signals – either in phase or in antiphase with one another – were uploaded, one to each AWG. These were written to span the full voltage range of the output.

We cycle between three measurements

1.	Net Gain	in-phase sinusoid, both electrodes on	Optical
2.	Gain Balance	in-phase sinusoid, electrodes individually	Electrical
3.	Inverted Balance	antiphase sinusoid, both on	Electrical

Where "Net Gain" sets the overall amplitude, which should be on the cusp of turning over when the phase spans the full 2π of the interferometer; "Gain Balance" is used to adjust the relative gain of the two electrodes, such that individually their amplitudes are balanced – i.e. that the two arms are symmetric by minimising their *difference* signal; and "Inverted Balance" seeks to minimise differences in the sign dependent response of the electrodes, achieved by minimising the *sum* signal between the two traces.

Strictly all of these can be achieved in the optical domain, however the electrode signal yields a much more favourable SNR; and is independent of the bias drift which could occur over the timescale of optimisation.

The three of these were "walked" until no changes (at the limit of the resolution of the power supply) made a noticeable difference.

In future the power supplies could be incorporated into the control libraries and we could conceivably perform a software optimisation of these – and real-time compensation of the bias voltage.

APPENDIX B

TEMPORAL ELECTRIC–FIELD AMPLITUDE AND PHASE: ALGORITHMIC RECONSTRUCTION VIA TIME–DOMAIN INTERFEROMETRY

In this appendix I present a technique we have developed for the reconstruction of the temporal amplitude and phase of pulses in order to facilitate the characterisaton of pulses interfaced with a quantum memory in order to attain the desired operation. Whether this is to achieve optimal interaction [139] or for the manipulaton of temporal modes of light [185] we seek the means to both generate and verify the complex spectra of these pulses.

There exists a well developed family of pulse characterisation techniques for ultrafast pulses, typically less than a picosecond; however, these methods (e.g. Frequency-Resolved Optical Gating: FROG [226], Spectral Phase Interferometry for Direct Electric-field Reconstruction: SPIDER [227], Mutual Interferometric Characterisation of two Electric-fields: MICE [228] etc.) are all based upon spectral measurement. These techniques, therefore, are not applicable to long pulses – such as those generally compatible with quantum memories – due to their narrow spectra.

The proposal for the scheme is due to B.B., and I then wrote the algorithm and initially built the interferometer. I was assisted in the lab by Sang Park (an M.Phys. student whom I co-supervised) who assisted the maintenance of the experiment and testing of the algorithm. The data presented herein were collected by myself.

B.1 Pulse Shaping Infrastructure

The heart of the pulse carving infrastructure is a fibre Mach-Zehnder interferometer in which the light is split into two paths. Each arm passes through an electro-optic crystal, in which the phase acquired can be arbitrarily controlled by the voltage applied to the electrodes V_i, for $i \in \{1, 2\}$.

The phase accrued in each path is then proportional to the voltage across each arm:

$$\phi_1 = (V_1 - V_{\text{bias}})/\alpha \tag{B.1}$$

$$\phi_2 = (V_{\text{bias}} - V_2)/\alpha \tag{B.2}$$

for some constant α. Then, for an input field $E_{\text{in}}(t)$, with 50/50 beamsplitters in the Mach-Zehnder,

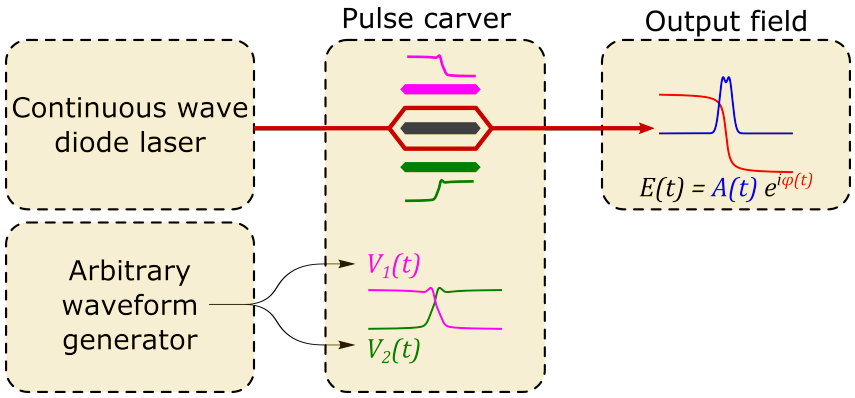


Figure B.1: Fibre Mach-Zehnder enabling abitrary amplitude/phase to be carved onto the output feld by modulating the electrode voltages on each arm of the interferometer.

the output is given by

$$\begin{aligned} E_{\text{out}}(t) &= \frac{\mathrm{i}E_{\text{in}}(t)}{2} \times \left(\mathrm{e}^{\mathrm{i}\phi_1} + \mathrm{e}^{\mathrm{i}\phi_2}\right) \\ &= \mathrm{i}E_{\text{in}}(t) \times \mathrm{e}^{\mathrm{i}\left(\frac{\phi_1+\phi_2}{2}\right)} \cos\left[\frac{\phi_1 - \phi_2}{2}\right] \end{aligned} \tag{B.3}$$

We therefore have arbitrary control of the amplitude modulation and phase (A and ϕ, respectively) of the output, namely

$$\begin{aligned} A &= \cos\left[\frac{\phi_1 - \phi_2}{2}\right] \\ &= \cos\left[\frac{V_1 + V_2 - 2V_{\text{bias}}}{2\alpha}\right] \end{aligned} \tag{B.4}$$

$$\begin{aligned} \phi &= \frac{\phi_1 + \phi_2}{2} \\ &= \frac{V_1 - V_2}{2\alpha} \end{aligned} \tag{B.5}$$

We have the freedom to set the bias voltage, V_{bias}, such that when no voltages are applied to the electrodes, $V_i = 0$, no light is transmitted, i.e. we then have:

$$A(t) = \sin\left[\frac{V_1(t) + V_2(t)}{2\alpha}\right] \tag{B.6}$$

$$\phi(t) = \frac{V_1(t) - V_2(t)}{2\alpha} \tag{B.7}$$

We can then calculate the voltage signals that must be applied to each electrode for an arbitrary output pulse profile via:

$$V_1(t)/\alpha = \arcsin\left[A(t)\right] + \phi(t) \tag{B.8}$$

$$V_2(t)/\alpha = \arcsin\left[A(t)\right] - \phi(t) \tag{B.9}$$

B.2 TEA+PARTI: Basic Idea

Here we introduce Temporal Electric-field Amplitude + Phase: Algorithmic Reconstruction via Time-domain Interferometry (TEA+PARTI). This constitutes a *single shot* method to fully characterise

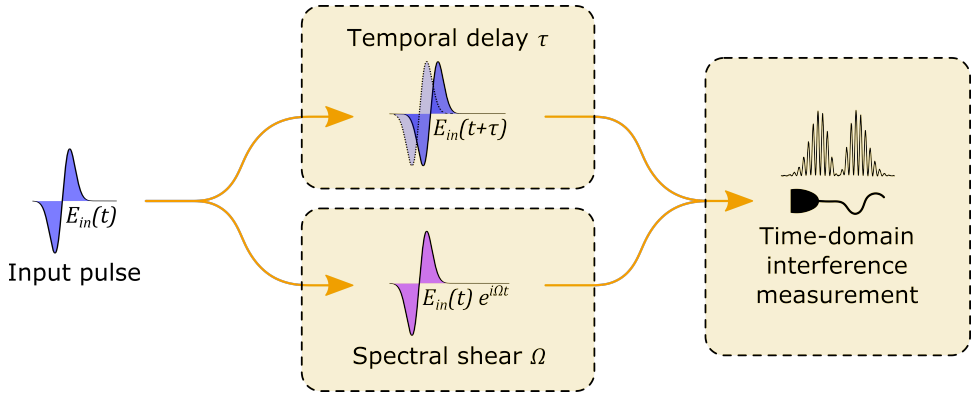


Figure B.2: Schematic for the interference measurement in TEA+PARTI.

temporal amplitude and phase of electric field.

Consider an incident electric field, with temporal envelope:

$$E_{\text{in}}(t) = A(t)\mathrm{e}^{\mathrm{i}\phi(t)} \tag{B.10}$$

$$\tilde{E}_{\text{in}}(\omega) = \mathcal{F}\left[E_{\text{in}}(t)\right] \tag{B.11}$$

for which we wish to reconstruct the temporal amplitude ($A(t)$) and phase ($\phi(t)$). In order to achieve this, we must therefore be sensitive to intensity ($\propto A^2$) and still access the spectral information ($\phi(t)$).

Analogously to SPIDER [227] developed in 1998 for the characterisation of ultrafast pulses we perform self referenced intererometry between a *temporally* delayed and *spectrally* sheared copy of the sample pulse, as depicted schematically in figure B.2. In contrast to SPIDER here we record the interference fringes in the time domain as opposed to performing a spectral measurement. This means that

In the time domain, this may be written

$$\begin{aligned} E_{\text{del}}(t) &= \alpha E_{\text{in}}(t+T) \\ &= \alpha A(t+T)\mathrm{e}^{\mathrm{i}\phi(t+T)}, \end{aligned} \tag{B.12}$$

$$\begin{aligned} E_{\text{shr}}(t) &= \beta E_{\text{in}}(t)\mathrm{e}^{\mathrm{i}\Omega t} \\ &= \beta A(t)\mathrm{e}^{\mathrm{i}[\phi(t)+\Omega t]}. \end{aligned} \tag{B.13}$$

Upon interfering the two copies and recording the output intensity as a function of time, this is proportional to:

$$\begin{aligned} S(t) &= \left|E_{\text{shr}}(t) + E_{\text{del}}(t)\right|^2 \\ &= \left\{\alpha^2 A^2(t+T) + \beta^2 A^2(t)\right\} \\ &\quad +\alpha\beta A(t)A(t+T)\left\{\mathrm{e}^{\mathrm{i}[\phi(t)-\phi(t+T)+\Omega t]} + \text{c.c}\right\} \\ &= \left\{\alpha^2 A^2(t+T) + \beta^2 A^2(t)\right\} \\ &\quad +2\alpha\beta A(t)A(t+T)\cos\left[(\phi(t)-\phi(t+T)) + \Omega t\right] \end{aligned} \tag{B.14}$$

We thus observe an envelope similar to the initial intensity profile, modulated by fringes dependent on the spectral shear and temporal delay between the copies of the two pulses. The fringe frequency comes primarily from the spectral shear Ω, with modification to this frequency from the

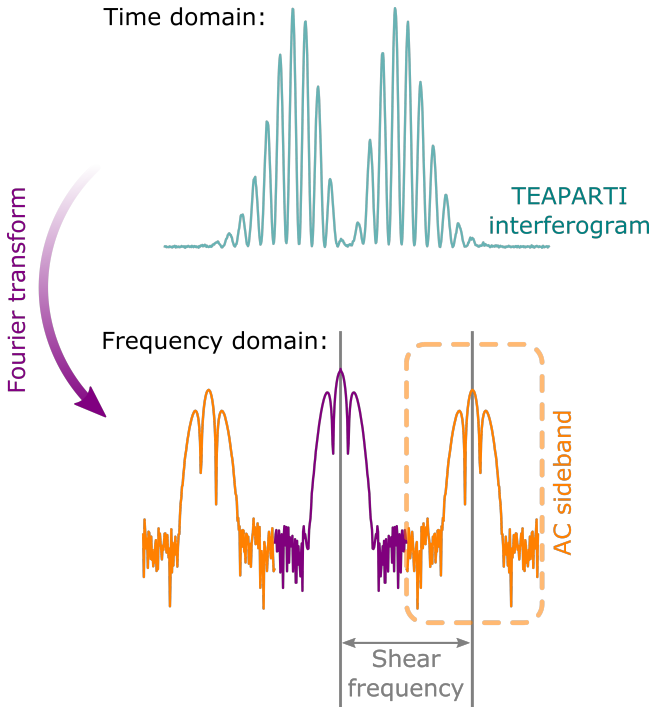


Figure B.3: fourier filtering of signal

term $(\phi(t) - \phi(t+T))$.

Transforming the measured trace into the frequency domain, by linearity of the Fourier transform, we can seperate out the constituent terms:

$$\begin{aligned}\tilde{S}(\omega) &= \mathcal{F}\left[\alpha A^2(t+T) + \beta^2 A^2(t)\right] \\ &\quad + \mathcal{F}\left[\alpha\beta A(t)A(t+T)\mathrm{e}^{-\mathrm{i}[\phi(t)-\phi(t+T)+\Omega t]}\right] \\ &\quad + \mathcal{F}\left[\alpha\beta A(t)A(t+T)\mathrm{e}^{+\mathrm{i}[\phi(t)-\phi(t+T)+\Omega t]}\right] \qquad \text{(B.15)} \\ &= \tilde{S}_0(\omega) + \tilde{S}_-(\omega) + \tilde{S}_+(\omega) \qquad \text{(B.16)}\end{aligned}$$

We thus see that we have side-band components at $\pm\Omega$. Provided that the spectral shear is larger than the spectral bandwidth of the signal, we may perform Fourier filtering of the signal and isolate the components $\tilde{S}_0(\omega)$, and the side-bands $\tilde{S}_\pm(\omega)$.

Transforming each of these *individually* back to the time domain, we have the DC component:

$$\begin{aligned}S_0(t) &= \mathcal{F}^{-1}\left[\tilde{S}_0(\omega)\right] \qquad \text{(B.17)} \\ &= \alpha A^2(t+T) + \beta A^2(t) \\ &= a^2(t) + b^2(t) \qquad \text{(B.18)}\end{aligned}$$

and the two AC components:

$$\begin{aligned}S_\pm(t) &= \mathcal{F}^{-1}\left[\tilde{S}_\pm(\omega)\right] \qquad \text{(B.19)} \\ &= \alpha\beta A(t)A(t+T)\mathrm{e}^{\pm\mathrm{i}[\phi(t)-\phi(t+T)+\Omega t]} \\ &= a(t)b(t)\mathrm{e}^{\pm\mathrm{i}\theta(t)} \qquad \text{(B.20)}\end{aligned}$$

where we have introduced the [*real*] variables $a(t)$ and $b(t)$ for the relative amplitudes, and $\theta(t)$ for

the interference phase for later convenience:

$$a(t) = \alpha A(t+T) \tag{B.21}$$
$$b(t) = \beta A(t) \tag{B.22}$$
$$\theta(t) = \phi(t) - \phi(t+T) + \Omega t \tag{B.23}$$

B.2.1 Amplitude Reconstruction

In terms of the quantities a and b, the magnitudes of the temporal components of the intensity spectrum are

$$|S_0| = a^2 + b^2 \tag{B.24}$$
$$|S_\pm| = ab \tag{B.25}$$

omitting the explicit time dependence for readability. Now if one considers constructing

$$\begin{aligned}
\frac{1}{2}\left[\left|\sqrt{|S_0|+2|S_\pm|}\right| \pm \left|\sqrt{|S_0|-2|S_\pm|}\right|\right] &= \frac{1}{2}\left[\left|\sqrt{(a^2+b^2)+2(ab)}\right| \pm \left|\sqrt{(a^2+b^2)-2(ab)}\right|\right] \\
\Rightarrow \frac{1}{2}\left[\left|\sqrt{|S_0|+2|S_\pm|}\right| + \left|\sqrt{|S_0|-2|S_\pm|}\right|\right] &= \frac{1}{2}\left[\left|\sqrt{(a+b)^2}\right| + \left|\sqrt{(\pm[a-b])^2}\right|\right] \\
&= \frac{1}{2}\left[|a+b| + |a-b|\right] \\
&= \begin{cases} a & \text{where } a > b, \\ b & \text{where } a < b, \\ 0 & \text{otherwise.} \end{cases} \qquad \text{(B.26)} \\
\Rightarrow \frac{1}{2}\left[\left|\sqrt{|S_0|+2|S_\pm|}\right| - \left|\sqrt{|S_0|-2|S_\pm|}\right|\right] &= \begin{cases} b & \text{where } a > b, \\ a & \text{where } a < b, \\ 0 & \text{otherwise.} \end{cases} \qquad \text{(B.27)}
\end{aligned}$$

we may directly access the amplitude information of the input pulse.[1] In practice, this may be done by evaluating the square root with both signs.

B.2.2 Phase Reconstruction

Looking at just one of the isolated side-band components, in the time domain:

$$S_\pm(t) = a(t)b(t)\mathrm{e}^{\pm \mathrm{i}\theta(t)} \tag{B.28}$$

[1] We may constrain the amplitudes a, b to be positive without loss of generality.

from which we may extract its phase, $\theta(t)$. If the applied shear can be accounted for, let $\Theta(t) = \Omega t - \theta(t)$, then we may evaluate the quantity:

$$\begin{aligned} \varphi(t + T/2) &= \frac{\Theta(t)}{T} && \text{(B.29)} \\ &= \frac{\phi(t+T) - \phi(t)}{T} \\ &\approx \left.\frac{\partial \phi}{\partial t}\right|_{t+T/2} && \text{(B.30)} \end{aligned}$$

which we may integrate directly, and obtain the original phase of the signal (up to a global phase):

$$\begin{aligned} \phi(t) &= \phi_0 + \int_{-\infty}^{t} \mathrm{d}t' \left.\frac{\partial \phi}{\partial t}\right|_{t'} \\ &\approx \phi_0 + \int_{-\infty}^{t} \mathrm{d}t' \varphi(t') && \text{(B.31)} \end{aligned}$$

Or, we can"concatenate" the phases from some chosen value $\phi(t_0)$, typically set to 0, then:

$$\begin{aligned} \phi(t_0 + T) &= \phi(t_0) + T \left.\frac{\partial \phi}{\partial t}\right|_{t_0} \\ &= \phi(t_0) + T \left[\frac{\phi(t_0 + T) - \phi(t_0)}{T}\right] \\ &= \phi(t_0) + \Theta(t_0) && \text{(B.32)} \end{aligned}$$

and so forth for subsequent multiples of T, generalising to

$$\phi(t_0 + NT) = \phi(t_0) + \sum_{n=0}^{N-1} \Theta(t_0 + nT) \qquad \text{(B.33)}$$

Note that this procedure samples the temporal phase at intervals of the temporal delay: so, for a pulse with spectral support Δ, the concatenation is exact provided that the delay satisfies $T \leq 2\pi/\Delta$.

B.2.2(a) Phase Stability

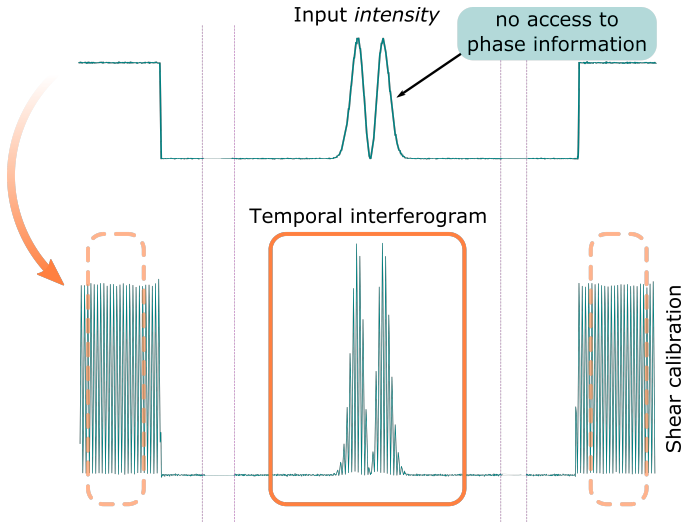
Note, in traditional SPIDER, the calibration removing the contribution from the delay phase is achieved by a repeat measurement *with zero delay*, or in the temporal case with zero shear ($\Omega = 0$). However, noting that, if the interferometer is not phase stable from shot to shot, the fringes of the recorded trace are washed out.

This may be modelled by expressing the fringe modulation as $\cos[\theta(t) + \epsilon_i]$, where for each repetition i the additional phase contribution ϵ_i has been included to encode perturbations on the differential phase between each arm of the interferometer, and ϵ_i is approximately constant within a single measumerent but $\epsilon_i \neq \epsilon_{j\neq i}$. i.e. $\langle \epsilon_i \rangle = 0$, but $\langle \epsilon_i^2 \rangle \sim \pi^2$. This means that the shear phase cannot be calibrated in a seperate measurement which would give

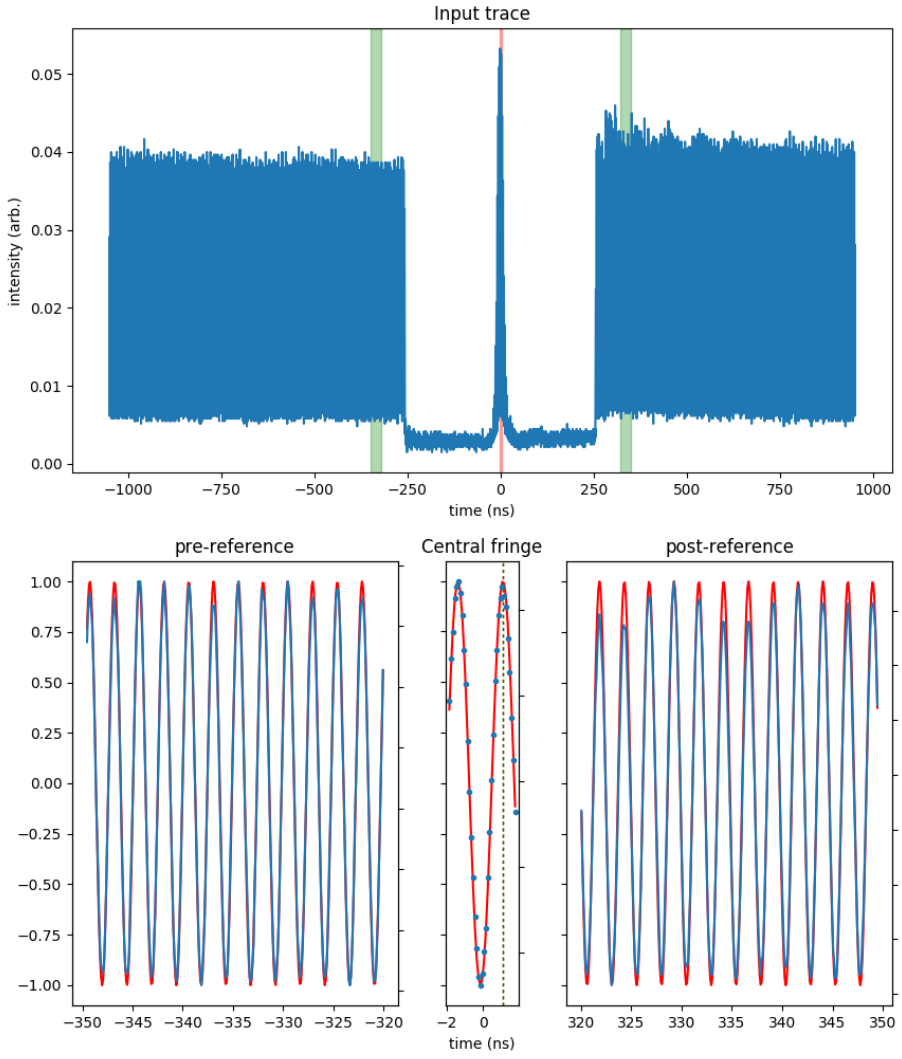
$$\theta_i(t) = \phi(t) - \phi(t + T) + \Omega t + \epsilon_i \qquad \text{(B.34)}$$

$$\theta_j^{T=0}(t) = \Omega t + \epsilon_j \qquad \text{(B.35)}$$

so even through obtaining a good measurement of $\langle \theta_i^{(0)} \rangle = \Omega t$, we would still have an additional shear. Thus, the shear must be measured for each measurement *in situ*.



(a) Shear calibration via CW wings



(b) Shear calibration in practice

Figure B.4: Calibration of shear frequency by allowing CW window before/after target pulse. (a) Operation principle, (b) implementation on real experimental data.

Skirting around this issue for now, let us simply consider that the delay and shear are taken to be calibrated and known (*cal.*), to some precision, with a small deviation from the true value (*tru.*) and qualify the consequences of such:

$$\Omega^{\text{cal.}} = \Omega^{\text{tru.}} + \Delta' \tag{B.36}$$

$$T^{\text{cal.}} = T^{\text{tru.}} + \delta' \tag{B.37}$$

Back of the envelope calculations for the corrections we have are of the form

$$\varphi(t) = \left[1 - \epsilon^{(1)}(\delta')\right] \left.\frac{\partial \phi}{\partial t}\right|_t + \epsilon^{(2)}(\Delta', \delta')t \tag{B.38}$$

With this in mind, if we attempt to integrate φ to find the phase:

$$\begin{aligned}\int_{-\infty}^{t} \mathrm{d}t' \varphi(t') &= \int_{-\infty}^{t} \mathrm{d}t' \left[(1 - \epsilon^{(1)}) \left.\frac{\partial \phi}{\partial t}\right|_{t'} + \epsilon^{(2)} t'\right] \\ &= \left[1 - \epsilon^{(1)}\right] \phi(t) + \epsilon^{(2)} t^2\end{aligned} \tag{B.39}$$

we see that uncertainty in the delay T, can introduce a prefactor scaling to the extracted phase. Similarly, uncertainty in the shear phase, Ω, therefore becomes manifest as an additional quadratic [chirp] component.

B.3 TEA+PARTI in Practice

We consider here the case that the delay calibration has been performed, and the only input we record is that of the interferogram - an illustrative example which is shown in figure B.4, where we include a period of CW transmission preceding and following the target pulse to calibrate the applied shear and carrier-envelope phase (CEP) in-situ. The delay was calibrated prior to each experimental run by recording reference traces passing through pulses in each arm of the interferometer independently. This is discussed further in section B.3.2.

Before that; as a reminder, let us reiterate the requirements of the shear and delay, for a pulse with *compact temporal support*[2] of T:

	Lower limit	Upper limit
shear, Ω	$\geq 2\pi/T$	
delay, T		$\leq T/2\pi$

B.3.1 Experimental Setup

The experimental setup for the TEA+PARTI demonstration is shown in figure B.5, in whiche the approximate interferometer parameters determined from the experiment were a spectral shear $\Omega = 2\pi \times 406\,\text{MHz}$ (double passed ~ 200 MHz AOM Crystal Technology) and temporal delay $T = -1.0$ ns. These data were obtained using the same infrastructure as in section 7.3.1, where the intensity reference pick off used in the main experiment used to monitor the optical output, was passed through the TEA+PARTI interferometer with the shear arm blocked and all intensity passing along the delay path.

[2] the time window within which the pulse energy is contained.

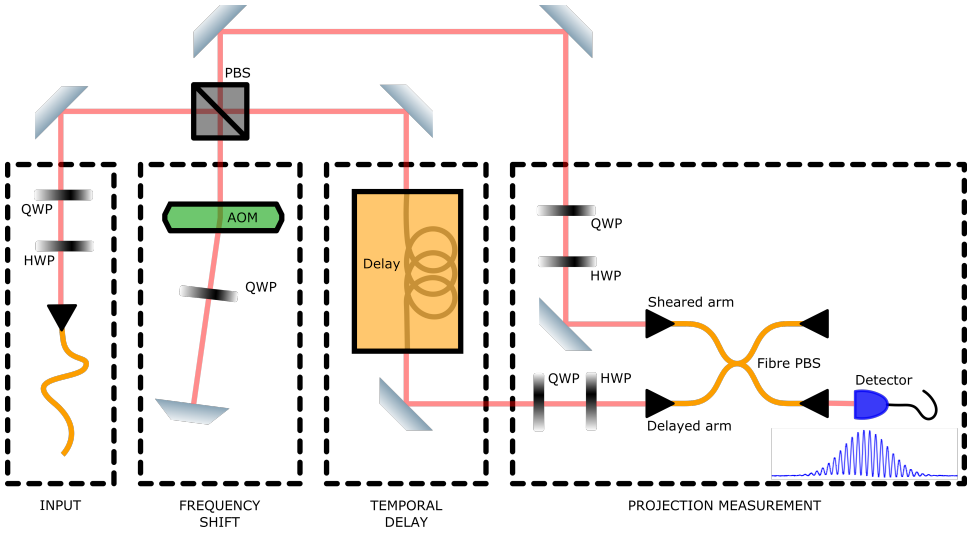


Figure B.5: Experimental setup of the TEA+PARTI interferometer.

B.3.2 Proper Calibration

Following from the previous discussion applying TEAPARTI to simulated data, we now move towards implementation in a real experimental situation. It is clear that proper calibration of the shear and delay are needed to accurately characterise the incident signal.

Preliminary measurements confirm that the CW interference, which by inspection of equation B.14 has the form:

$$S_{\mathrm{cw}}(t) = a + b\cos\left[\Omega t + t_0\right]. \tag{B.40}$$

In this preliminary implementation of the experiment, this is phase stable with the reference over the duration of one measurement cycle. Note that, to improve the stability of the experiment, the interferometer is enclosed in a box. In this expression, a and b are time constant during the CW interference, and we introduce t_0 to encode the CEP of the applied shear.[3]

Fitting the CW interference in the wings gives us an estimate of the shear Ω. This directly gives us the shear Ω. Then, in order to extract the CEP of the interferogram, Ωt_0, such that the shear phase can be correctly be removed, we allow a period of CW transmission prior to- and following- the pulse of interest, see figure B.4. Note that these must be seperated enough, in time, such that there is no transient response from this which may affect the pulse of interest. This is necessary, since we lack the stability to extract this offset once from the REF signal, as this will vary shot-to-shot.

The temporal delay, T, is the only parameter obtained in a seperate measurement. For this, each arm of the interferometer is blocked in turn and we find the position of the pulses with respect to the AWG trigger, refer to figure B.5

Having obtained measurements for the parameters Ω, T and t_0, we then normalise frequency/time variables with respect the applied shear. Defining the shear period $T_\Omega = 2\pi/\Omega$, we can now work in

[3] It is this t_0 which gives rise to the ϵ_i of the previous section, and varies from shot to shot.

terms of the dimensionless variables:

$$\Omega \to 1 = \Omega \times T_\Omega / 2\pi \tag{B.41}$$

$$T \to \delta = T / T_\Omega \tag{B.42}$$

$$t \to \tau = (t - t_0) / T_\Omega \tag{B.43}$$

$$t_0 \to 0 \tag{B.44}$$

The shear phase to be removed can then be rewritten:

$$\begin{aligned} \exp\left[\mathrm{i}\Omega\,(t - t_0)\,\right] &= \exp\left[2\pi\mathrm{i}\,\left(\Omega \times \frac{T_\Omega}{2\pi}\right)\,\left(\frac{t - t_0}{T_\Omega}\right)\,\right] \\ &\to \exp\left[2\pi\mathrm{i}\,\tau\right] \end{aligned} \tag{B.45}$$

which is zeroed upon the central fringe of the cw interference trace. When we then come to construct the phase, we can simply obtain Θ from:

$$\mathrm{e}^{\mathrm{i}\theta(\tau)} / \mathrm{e}^{2\pi\mathrm{i}\tau} = \mathrm{e}^{-\mathrm{i}\Theta(\tau)} \tag{B.46}$$

and the concatenation follows as:

$$\begin{aligned} \text{Let} \quad & \Theta(\tau) = \phi(\tau + \delta) - \phi(\tau). \\ \text{Set} \quad & \phi(0) = 0, \\ \text{then} \quad & \phi(\delta) = \Theta(0). \end{aligned} \tag{B.47}$$

$$\cdots$$

$$\text{So}: \quad \mathrm{e}^{\mathrm{i}\phi(N\delta)} = \prod_{n=0}^{N-1} \mathrm{e}^{\mathrm{i}\Theta(n\delta)} \tag{B.48}$$

Some examples of reconstructed pulses are shown in figure B.6. We plot the target, inferred and reconstructed fields for comparison, along with the amplitude obtained by the single arm intensity transmission in a separate measurement. Hereagain the the bias drift of the FMZI may be problematic as this information is not captured by the measurement of the electrode voltages.

B.4 Applications

Whilst the pulses considered here are much longer than those used in chapter 7, we may still use the results here to obtain an intuition for the validity of inferring the optical output field from the electrode voltages.

In most cases we obtain good agreement between inferred fields and those reconstructed using TEA+PARTI, typically obtaining an overlap $\gtrsim 70\%$ as computed by the numerical integrals of the complex fields.

In some cases this was as low as 20% for the COMP1 and 50% for the MUB11 – although, for those we note there was also the largest mismatch between the observed optical fields (figure 7.3) and the inferred signals (figure 7.1), observed for the short pulses.

Performing simulations with comparable SNR – where random field noise (amplitude and phase) is added to the input; each arm of the interferometer; and to the trigger signal and AOM driver. Then

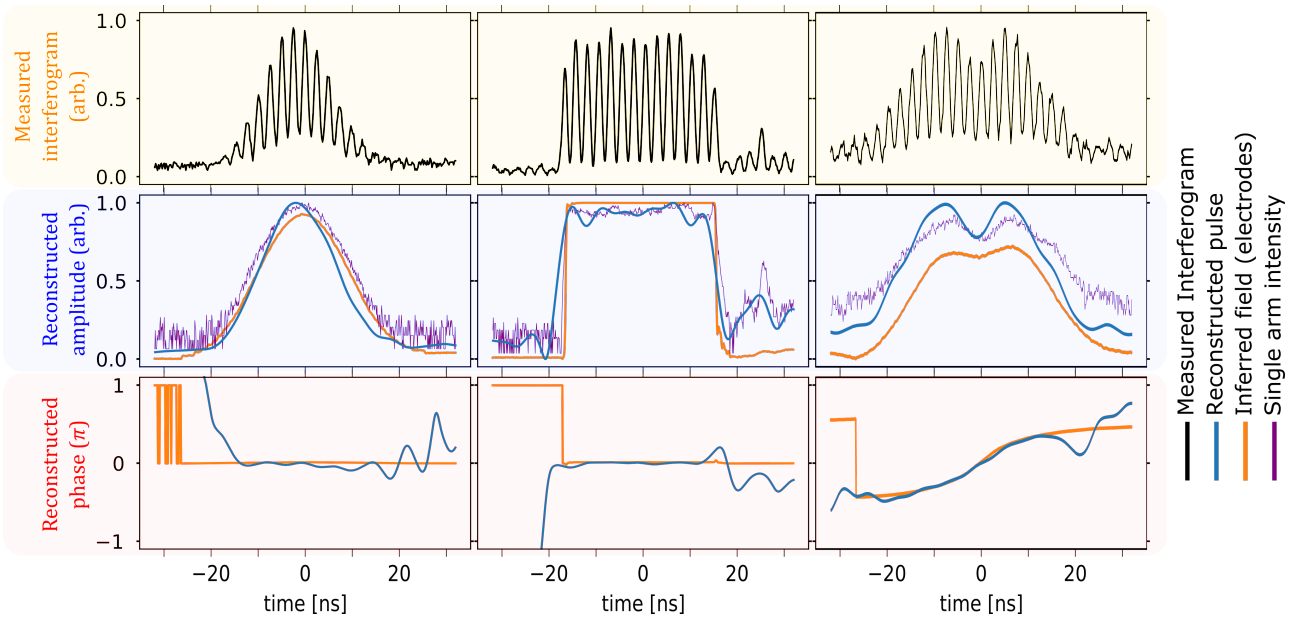


Figure B.6: Example output pulse profiles, experimentally reconstructed using TEA+PARTI – from a single measurement of the interferogram we obtain the reconstructed pulse, which is compared with the field profile obtained from the electrodes, as well as the amplitude obtained from the single arm intensity.

using the *same algorithm* as used in experiment to extract the delay and shear from noisy simulated data, we should in principle still be able to extract the amplitude and phase information. An example of such simulation is shown applied to the COMP1 state is shown in figure B.7.

In this example, we observe there is a slight [negative] chirp contribution, and [overestimate] in magnitude in the phase reconstruction. From the simple error propagation these may be expected from equation B.39, since the shear and delay are not extracted perfectly (the calibrations are off by approximately -1 MHz and -0.1 ns respectively).

Even in this case, the algorithm performs robustly – and handles the zero in intensity/phase flip reliably[4] – and so either there is some additional noise process or physics not captured in this model, which contributes to the reconstruction on experimental data.

One such example of "physics not included" is the bias drift: it is likely that this remains a significant obstacle – and drifts on the minute timescale of a few measurements – in between calibrating the delay and checking the bias and measuring the interferogram. This information is not captured by the measurement of the electrode voltages, and so accounts for at least some of the discrepancy between the reconstructed and inferred signals.

A brief attempt was made to perform a fit minimising the difference between the reconstructed field and that inferred from the electrode voltages, as the (i) bias voltage, (ii) shear and (iii) delay mismatch were varied, however with this parametrisation fitting routines did not perform reliably or converge to a physically sensible result. – this may be worth pursuing in order to verify whether or not this description captures sufficient physics, or if there are additional factors we must include.

When characterising the set up – it was noted by S.P. that in reconstructions on Gaussian pulses with simple phase profiles (e.g. linear ramp or quadratic) that erroneous phases were seen in the reconstructions that depended on the form of the phase profile writen; and it was concluded that this

[4] Whilst zeros in intensity may be well captured in more recent interferometric techniques FROG or MICE, which typically involve combining multiple measurements scanning some paramater, phase flips can be problematic in traditional SPIDER.

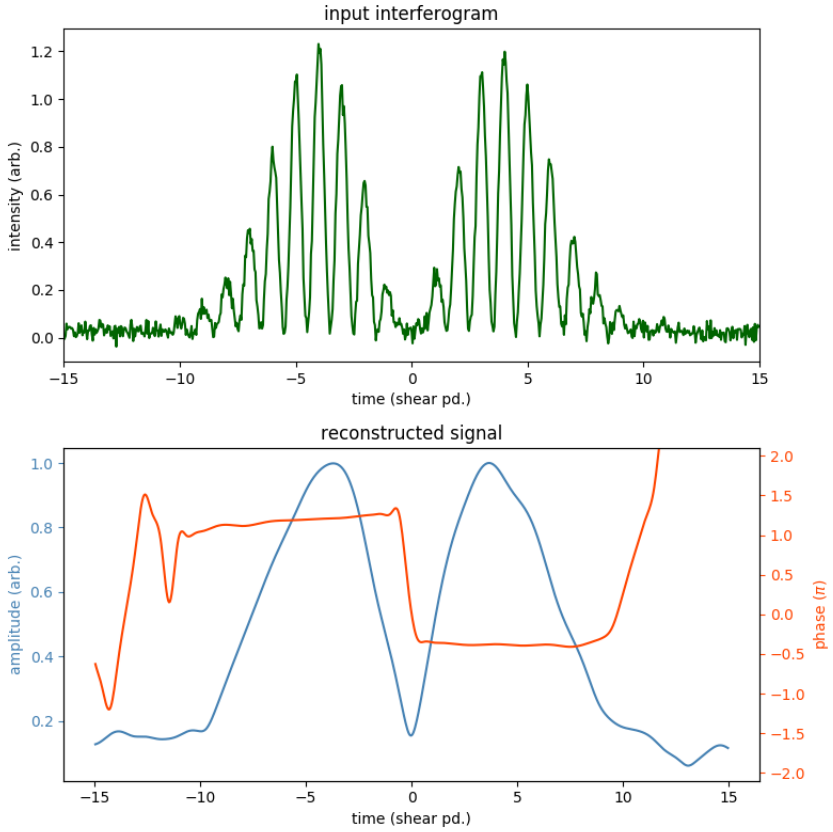


Figure B.7: Typical simulation data with realistic noise (described in the main text) applied to the COMP1 state, run through the same reconstruction algorithm.

was most likely due to imperfections in the EOM – for instance that the change in refractive index of the waveguide is not perfectly linear with the applied voltage [229].

This observation suggests that – once confident of the robustness of TEA+PARTI – implementing an adaptive algorithm to obtain the target field based upon direct reconstruction of the resulting field[5] may be an appropriate way forward.

B.4.1 Interfacing with a Memory

As discussed in the body of this thesis, the two principal applications we have in mind for shaped pulses, and where this protocol would constitute a valuable technique:

- exploiting the single-mode nature of the Raman memory: performing selective operations determined by the temporal mode of the signal/control fields
- compensation for the AC Stark shift

The components for this set up were chosen with the cavity raman memory discussed in the part I of this thesis, which is compatible with "longer pulses" but not immediately applicable to the case studied in part II.

However in future versions of the memory expeiment in which longer duration control pulses may be used (if also carved using this system, then amplified)[6]

[5] Or, optimising on some other figure of merit, and characterising the optimal field retrospectively.

[6] Or if the AOM is replaced to implement a single pass shear $\gtrsim 300$ MHz, although in this regime the amplifier ringing observed limit our ability to carve the desired pulses faithfully.

B.4.2 Characterisation of Single Photons

An appealing future application of this method would be performing the full temporal reconstruction of single photons. However, this introduces further difficulties.

In this case we would consider carving pulses which could then be amplified by a TA and used to pump a PDC source, and characterising the resulting temporal mode of the daughter photons.

- no longer single shot - need to acquire statistics
  - phase stability of interferometer for long enough durations
  - eom bias drift needs to be actively stabilised
- low temporal jitter photon detection to resolve fringes
  - Nanowires, low jitter trigger (e.g. optical) needed
- phase stable shear frequency
  - whilst the AOM was stable as far as we could measure, the clock for the AWG would need to be locked with the AOM driver to constrain the CEP in each shot.

B.5 Conclusions

The results in this appendix still require further development before being deployed as a robust characterisation technique. This was not pursued further at the time as the equipment available meant that this couldn't be used *directly* for the pulses used in the memory experiments in the body of this thesis; however the initial results are a promising step towards enabling the full characterisation of the complex spectrum of ns duration pulses compatible with quantum memories.

With some improvements to the experimental stability, and some further characterisation this method could be an invaluable tool.

PS

As we conclude in the main text, the future development of the experiment will include using a tapered amplifier (TA) to amplify the carved pulses for use as the control field in the memory interaction, and could therefore operate with pulses of compatible bandwidth.

This would also facilitate further characterisation of TEA+PARTI – with, say, a much improved SNR . The amplified pulses would have power to spare to pick off before the interferometer and could conceivably implement a subroutine interleaved with the main measurements to set the bias (by minimising transmission when the electrodes are off), and monitoring the intensity profile *simultaneously* as a means of error checking.

Including the shutters (such as those in appendix A.2) to speed up the duty cycle between measurement configurations and overconstraining the reconstruction could be beneficial.

Having said that, the TA itself is likely to induce some pulse distortion so comparison with the electrode voltages would no longer be such a useful metric.

I anticipate that rebuilding the experiment with improved stability and more real time monitoring should enable TEA+PARTI to become a useful tool for the future work interfacing with temporal modes in a spectral bandwidth regime otherwise inaccessible by standard pulse characterisation techniques.

APPENDIX C

OPTIMISATION OF AWG ELECTRODE VOLTAGES

In the main text it was noted that the preparation of some target optical signal could only be achieved with limited fidelity, given the restricted sampling rate available for the AWG outputs – particularly whilst synchronised with the master laser – and the response function of electrodes.

In addition to developing the means to *directly* characterise optical outputs via TEA+PARTI, a complementary approach is to vary the AWG signals to optimise upon the output. As seen in appendix B, measurement of the electrodes gives us some confidence in estimating the true optical field, nevertheless in figures 7.1 and 7.3 we see that (i) the amplifiers do not perform ideally; and (ii) the inferred optical signals do not fully agree with those measured.

As a first attempt to address this, we consider methods that may be used to compensate these imperfections – in the first instance I performed some preliminary investigations based upon optimising upon the measured amplifier outputs in order to improve the fidelity with which a given signal can be prepared. Ultimately the cost-function could be evaluated at later stages in the experiment, for instance with direct reconstruction techniques, or optimisation of memory performance.

Whilst not directly used in any experiment presented in the main thesis, this side project was principally aimed to utilise (and test!) the control infrastructure I have developed for the experiment, and demonstrate the potential for further automation and on-the-fly optimisation of the experiment in the future, rather than development of a suitable search algorithm.

C.1 Downhill Simplex Search

C.1.1 Search Algorithm

The downhill simplex search, or Nelder-Mead algorithm is a traditional method for direct, *brute force*, optimisation. It is known that whilst Nelder-Mead methods perform well in low dimension, it can become rather inefficient in high-dimensional problems; nevertheless, it is intuitively simple to understand, and straightforward to implement.

We seek to solve some unconstrained optimisation – parametrised by finding $\boldsymbol{x}$ that minimises some cost-function $f(\boldsymbol{x})$, where the objective function f maps some n dimensional problem to a scalar cost.

Within the n dimensional space, a *simplex* is defined by $n+1$ vertices, $\{\boldsymbol{x}_i\}_{i=0}^{n}$, and encloses some region – the convex hull. At each vertex we may evaluate $f(\boldsymbol{x}_i)$, and in each stage of the algorithm

we generate a new simplex, by obtaining a new estimate of the optimal coordinate $\boldsymbol{x}_{\text{opt}}$ based upon the vertex evaluations, and discard the worst.

The algorithm is based upon four elementary operations:

1.	Reflection:	α
2.	Expansion:	β
3.	Contraction:	γ
4.	Shrink:	δ

At some iteration, j, of the algorithm, we re-order the vertices, i, according to their cost function

$$f\left(\boldsymbol{x}_0^{[j]}\right) \leq f\left(\boldsymbol{x}_1^{[j]}\right) \leq \ldots \leq f\left(\boldsymbol{x}_n^{[j]}\right)$$

from which we define the centroid based upon all but the worst vertices:

$$\bar{\boldsymbol{x}}^{[j]} = \frac{1}{n}\sum_{i=0}^{n-1} \boldsymbol{x}_i^{[j]}$$

and the vertices of new coordinates to evaluate, based upon the geometric operations (I suppress the iteration superscript):

$$\begin{aligned}
\text{Reflection} \quad & \boldsymbol{x}_r &&= \bar{\boldsymbol{x}} + \alpha\left(\bar{\boldsymbol{x}} - \boldsymbol{x}_n\right) \\
\text{Expansion} \quad & \boldsymbol{x}_e &&= \bar{\boldsymbol{x}} + \beta\left(\boldsymbol{x}_r - \bar{\boldsymbol{x}}\right) \\
\text{Out-contraction} \quad & \boldsymbol{x}_{oc} &&= \bar{\boldsymbol{x}} + \gamma\left(\boldsymbol{x}_r - \bar{\boldsymbol{x}}\right) \\
\text{In-contraction} \quad & \boldsymbol{x}_{ic} &&= \bar{\boldsymbol{x}} - \gamma\left(\boldsymbol{x}_r - \bar{\boldsymbol{x}}\right)
\end{aligned}$$

and finally, a "shrinking" operation, where all vertices $i = 1, ..., n$ are replaced, shrinking the simplex about the *best* vertex $\boldsymbol{x}_0$, by $\boldsymbol{x}_i^{[j]} = \boldsymbol{x}_0^{[j]} + \delta\left(\boldsymbol{x}_i^{[j]} - \boldsymbol{x}_0^{[j]}\right)$.

We adopt the parameters proposed in [230] to improve the performance of the Nelder-Mead method for larger n dimensional searches:

$$\alpha = 1, \qquad \beta = 1 + \frac{2}{n}, \qquad \gamma = 0.75 - \frac{1}{2n}, \qquad \delta = 1 - \frac{1}{n}.$$

The initial simplex is constructed from the first vertex at the vector parameterising our initial estimate, and the remaining n vertices by perturbing each coordinate by the unit vector along each coordinate axis in turn.[1]

Running for a fixed number of cycles, as the simplex evolves in each iteration, we monitor the cost of the best vertex. If this plateaus within some tolerance, we take our current best guess as a new "initial vertex" from which a new simplex is randomly generated as before – the lower figure of C.3 shows the best/wost costs at each iteration of the algorithm, where the spikes indicate each step at which a new simplex is generated.

C.1.2 Parameterisation of Search Problem

In this test implementation, we begin from our target electrode inputs as output from the AWGs sampled at 50 Gs/s. We define our simplex vertices as a vector of 32 coordinates distributed across the pulse where we displace the signal. We then interpolate between these points to define a smooth perturbation at the AWG sample rate, which is added to the target signal.

[1] Or, drawn from a random-normal distribution $\sigma = 1$

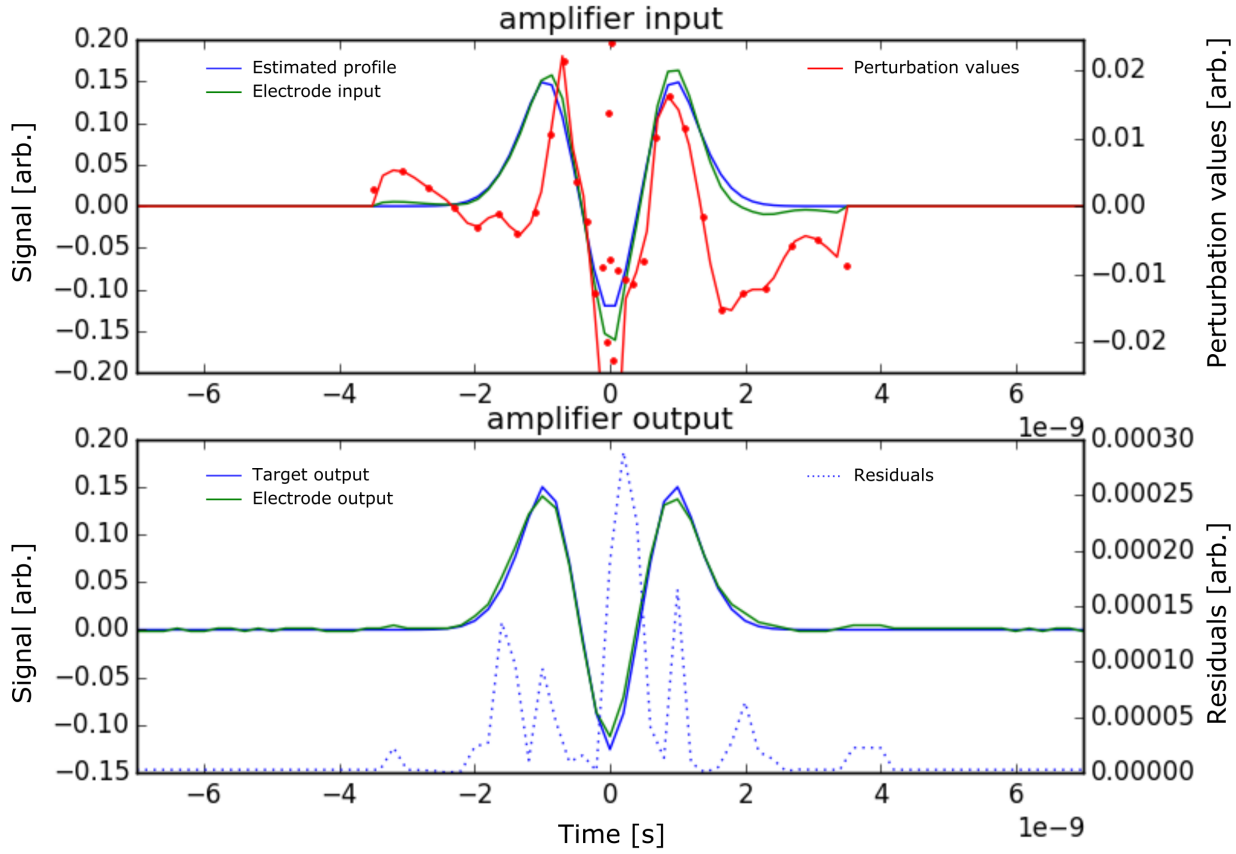


Figure C.1: Possible parametrisation of search vector perturbing target electrode voltage, described in the text.

This parametrisation is illustrated in figure C.1 shown for an intermediate step of the algorithm for one of the electrode signals for the generation of an HG2 pulse. The blue curve displays the target signal, and the green is the written (upper) and resulting (lower) amplifier signal, and the red points the perturbation points. The dashed blue line shows the residuals on the amplifier at the perturbation points.

C.1.3 Simplex Iteration

Figure C.2 shows the flow of an iteration of the algorithm. We highlight again that each vertex evaluation requires that the corresponding waveforms be evaluated, loaded to the AWGs and then run and measured by the 'scope. As such the algorithm has significant time overheads, with the evaluation of cost at each vertex taking a few seconds.

Typically each iteration will only involve the evaluation of 1-4 vertices, periodically the whole simplex will be regenerated – see figure C.3 showing an example of the evolution of the cost function – where the peaks in the lower figure are seen when the simplex is randomly regenerated.

C.1.4 Test

The limited response of the electrodes becomes more problematic for higher order functions, as temporally finer features become necessary. For this preliminary investigation I looked at optimisation

1. **Sort:** Index vertices in increasing cost order $f_i = f(\boldsymbol{x}_i)$.
2. **Reflection:** Evaluate cost f_r at reflected vertex $\boldsymbol{x}_r$
   - *if* ($f_r < f_0$)*:* go to 3...
   - *else if* ($f_0 \leq f_r < f_{n-1}$)*:* replace $\boldsymbol{x}_n$ with $\boldsymbol{x}_r$, return to 1...
   - *else if* ($f_{n-1} \leq f_r < f_n$)*:* go to 4...
   - *else* ($\Rightarrow f_n \leq f_r$)*:* go to 5...
3. **Expansion:** Evaluate cost f_e at expansion vertex $\boldsymbol{x}_e$
   - *if* ($f_e < f_r$)*:* replace $\boldsymbol{x}_n$ with $\boldsymbol{x}_e$, return to 1...
   - *else* ($\Rightarrow f_r \leq f_e$)*:* replace $\boldsymbol{x}_n$ with $\boldsymbol{x}_r$, return to 1...
4. **Out-Contraction:** Evaluate cost f_{oc} at outside contraction point $\boldsymbol{x}_{oc}$
   - *if* ($f_{oc} < f_r$)*:* replace $\boldsymbol{x}_n$ with $\boldsymbol{x}_e$, return to 1...
   - *else* ($\Rightarrow f_r \leq f_{oc}$)*:* go to 6...
5. **In-Contraction:** Evaluate cost f_{ic} at inside contraction point $\boldsymbol{x}_{ic}$
   - *if* ($f_{ic} < f_n$)*:* replace $\boldsymbol{x}_n$ with $\boldsymbol{x}_e$, return to 1...
   - *else* ($\Rightarrow f_r \leq f_{oc}$)*:* go to 6...
6. **Shrink simplex!**

Figure C.2: Iteration of traditional Nelder-Mead algorithm, summarised from [231].

for the HG1 and HG2 pulses, aiming to compensate for overshoot/damping observed in the signals.

Some typical results are shown in figures C.3 showing the evolution and C.4 showing some final results.

C.2 Future

Whilst the downhill simplex search is certainly not necessarily efficient, or even reliable, for high-dimensional searches (in this case we chose $d = 32$) this preliminary investigation demonstrates the feasibility of using the infrastructure I have developed for the optimisation of the preparation of shaped pulses.

For the purposes of the memory characterisation measurements in chapter 7, we don't need perfect preparation of the states, we just need to know what the optical fields are (see appendix B), for future measurements higher fidelity state preparation would be desirable – as such this sort of method could be applicable to optimising the AWG signals based upon a cost function derived from a direct reconstruction of the optical field (if the signals are compatible with TEA+PARTI) – in this case if a similar parametrisation approach were used, the dimensionality of the search space would double. The time overheads would significantly increase, and ensuring the long term stability of the experiment – in particular the bias drift of the FMZI – otherwise we would be attempting to optimise upon a dynamically changing landscape!

An alternative optimisation approach would be to directly optimise upon some other cost, for instance on memory efficiency – which could be automated integrating with the infrastructure discussed in A.3.1, and only evaluate the optical fields after some optimum has been found.

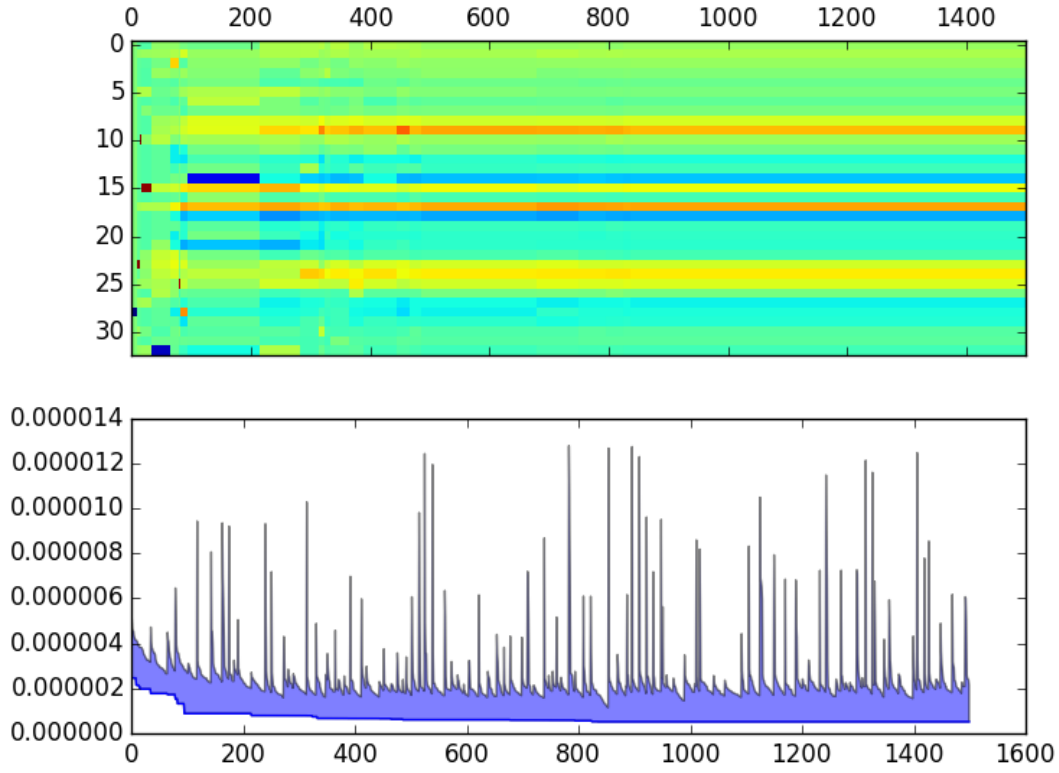
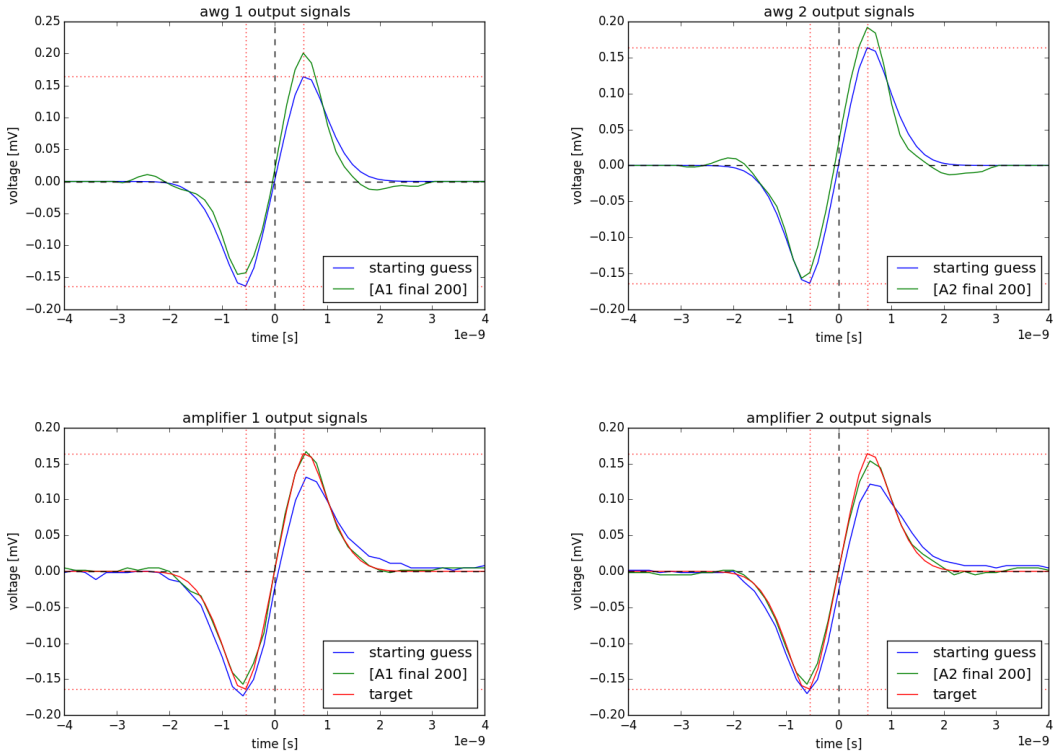
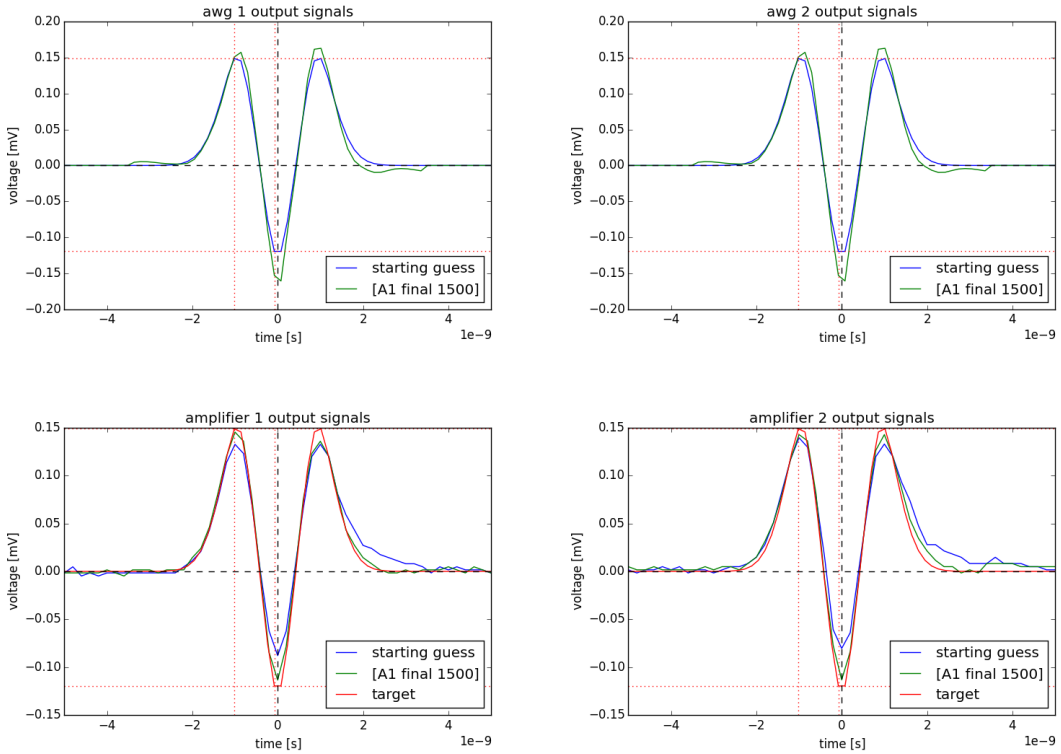


Figure C.3: Simplex evolution (top) colour plot of the coordinates of the simplex (bottom) evolution of the cost function – the shaded area covers the range of vertex costs across the whole simplex; the spikes are where the algorithm has stalled and a new, random simplex is generated based upon the current best guess. Note that the absolute value of the cost function depends on the choice of definition.



(a) HG1, 200 iterations



(b) HG2 1500 iterations

Figure C.4: Traces showing the target electrode signals, initial response and final response (a) for HG1 after 200 iterations and (b) HG2 after 1500 iterations. The blue curves show the unoptimised input/outputs, displaying overshoot and ringing; red curves display the desired response, with extrema highlighted by dashed lines; and green show the final input/output showing significantly improved agreement with the target output.

APPENDIX D

RECONSTRUCTION ALGORITHM

Quantum tomography describes the process of acquiring information about the state of a system from a series of measurements.

In the main text we seek to reconstruct the Green's function of the memory interaction over many attempts by successively probing the system with a range of different probe states and observing the "frequency" (or efficiency) with which they are stored by a particular memory system.

This is analogous to the methods in quantum state tomography where measurement of the quantum state of a system requires an ensemble of measurements to be applied to many copies of the same system (unlike the classical case, where one system is needed); and from the probabilistic results of the different measurements, we may build up a consistent description defining our state of knowledge of the system [232].

The problem then, is to find the system state which is the most likely, given the measurement outcomes we have observed.

D.1 Informationally Incomplete Maximum Likelihood Tomography

I summarise the derivation of the diluted maximum likelihood tomography method adapted from reference [207], using the proofs in [206] which is implemented in section 7.4 of the main text.

Given *many copies* of some unknown state ρ, upon each of which we may perform a set of measurements $\{\Pi_j\}$; we wish to reconstruct ρ based upon the observed measurement results. If we consider a set of such measurements, the probability of observing an outcome j is

$$p_j(\rho) = \mathrm{Tr}\left[\rho\Pi_j\right]. \tag{D.1}$$

Extended to an ensemble of N measurements, the *likelihood*, $\mathcal{L}$, of observing some set of results, given the underlying density matrix ρ can be expressed

$$\mathcal{L}(\rho) = \frac{N!}{\prod_k (n_k!)} \prod_j p_j(\rho)^{n_j} \tag{D.2}$$

$$\mathcal{L}(\rho)^{1/N} = a^{1/N} \prod_j p_j(\rho)^{f_j} \tag{D.3}$$

where a is some unimportant normalisation constant, and $f_j = n_j/N$ are the relative frequencies (i.e. $\sum_j f_j = 1$).

What we therefore seek to do is find an appropriate estimate of ρ which maximises $\mathcal{L}(\rho)$. Typically it is more tractable to consider the "log-likelihood":

$$\ln\left[\mathcal{L}(\rho)\right] = \sum_j f_j \ln\left[p_j(\rho)\right] + \text{const.} \tag{D.4}$$

and consider variations in our estimate of ρ yielding extremal likelihood.

This optimisation can be thought of multiple ways; for instance using the calculus of variations where we seek to extremise log-lkelihood over ρ. Consider for example the functional

$$F = \sum_j f_j \ln\left(\text{Tr}\left[A^\dagger A \Pi_j\right]\right) - \lambda \text{Tr}\left[A^\dagger A\right] \tag{D.5}$$

in which we have enforced the positivity of the density matrix by setting $\rho = A^\dagger A$, and the normalisation constrained by the Lagrange multiplier λ.

The extremal state is stationary under variation in $A \to A + \delta A$, and we require that $\delta F = 0$:

$$\delta F = \sum_j \frac{f_j}{p_j} \text{Tr}\left[\Pi_j A^\dagger \delta A\right] - \lambda \text{Tr}\left[A^\dagger \delta A\right] \tag{D.6}$$

$$\Rightarrow \quad \sum_j \frac{f_j}{p_j} \Pi_j A^\dagger = \lambda A^\dagger \tag{D.7}$$

With this in mind, we define:

$$\mathbb{R}(\rho) = \sum_i \frac{f_i \Pi_i}{\text{Tr}\left[\rho \Pi_i\right]} = \mathbb{1}_\rho \tag{D.8}$$

for some state ρ, and set of observed frequencies f_i. This may be interpereted as a closure relation, where $\mathbb{1}_\rho$ represents the identity operator on the support of an extremal density matrix (i.e. one which extremises the likelihood) ρ_e.

One derivation makes use of Jensen's inequality, namely that for some vectors $\{x_i, c_i\}$ with frequencies f_i, that we have

$$\prod_i \left(\frac{x_i}{c_i}\right)^{f_i} \leq \sum_i f_i \left(\frac{x_i}{c_i}\right)$$

$$\Rightarrow \quad \prod_i x_i^{f_i} \leq \left(\prod_j c_j^{f_j}\right) \sum_i f_i \left(\frac{x_i}{c_i}\right)$$

with equality iff $x_i/c_i = \text{const.}\, \forall i$. Applying this to the expression for the likelihood D.3,

$$\prod_j p_j(\rho)^{f_j} \leq \left(\prod_j c_j^{f_j}\right) \sum_i p_i \left(\frac{f_i}{c_i}\right)$$

now recalling $p_i = \text{Tr}\left[\rho \Pi_i\right]$, the sum on the right may be rewritten

$$\sum_i \text{Tr}\left[\rho \Pi_i\right] \left(\frac{f_i}{c_i}\right) = \text{Tr}\left[\rho \sum_i \left(\frac{f_i \Pi_i}{c_i}\right)\right] = \text{Tr}\left[\rho \mathbb{R}(\rho_\text{e})\right]$$

where we have set $c_i = \text{Tr}\left[\rho_\text{e} \Pi_i\right]$. This therefore defines our upper bound for the likelihood, and we

must have that

$$\mathrm{Tr}\left[\rho\mathbb{R}(\rho_{\mathrm{e}})\right] \leq \lambda \tag{D.9}$$

where λ is the largest eigenvalue of $\mathbb{R}(\rho_{\mathrm{e}})$, and an extremal state will correspond to the maximal eigenvalue. Furthermore, for an estimate of the state ρ which maximises the likelihood, then $x_i/c_i = \mathrm{Tr}\left[\rho\Pi_i\right]/\mathrm{Tr}\left[\rho_{\mathrm{e}}\Pi_i\right]$ must be stationary, hence we also reach equality in equation D.9.

When these are satisfied, we may therefore write $\mathbb{R}\rho = \rho$. In general it is not necessarily possible to solve this by analytical means; however, since both ρ and $\mathbb{R}$ are Hermitian, it follows that the relation

$$\mathbb{R}\rho\mathbb{R} = \rho \tag{D.10}$$

also holds, and both sides are now enforced to be positive semidefinite.

With no prior knowledge, starting from an *unknown* density matrix, our initial guess must be unbiased – for instance taking the maximally mixed state $\rho_0 = \mathbb{1}/\dim(\mathcal{H})$.

Applying equation D.10 directly, convergence to the true state is not guaranteed. Nevertheless, if ρ is iteratively *diluted* in sufficiently small steps, then this can be shown to strictly increase the likelihood at each step.

More explicitly, for some small positive dilution parameter ε, we define a modified propagator:

$$\tilde{\mathbb{R}}(\rho) = \frac{\mathbb{1} + \frac{\varepsilon}{2}\mathbb{R}(\rho)}{1 + \frac{\varepsilon}{2}}$$

and so at some iteration k of the dilution, the next estimate for ρ is evaluated using the current state $\rho^{(k)}$ and estimate of the propagator $\mathbb{R}^{(k)} = \mathbb{R}(\rho^{(k)})$ using:

$$\begin{aligned}\rho^{(k+1)} &= \tilde{\mathbb{R}}^{(k)}\rho^{(k)}\tilde{\mathbb{R}}^{(k)} \\ &\approx (1-\varepsilon)\left[\rho^{(k)} + \frac{\varepsilon}{2}\left(\mathbb{R}^{(k)}\rho^{(k)} + \rho^{(k)}\mathbb{R}^{(k)}\right) + \mathcal{O}\left(\varepsilon^2\right)\right] \\ &= (1-\varepsilon)\rho^{(k)} + \varepsilon\bar{\rho}^{(k)}\end{aligned}$$

i.e. constructed from a *convex mixture* of $\rho^{(k)}$ and $\bar{\rho}^{(k)}$, where we have defined

$$\bar{\rho}^{(k)} = \frac{1}{2}\left[\mathbb{R}^{(k)}\rho^{(k)} + \rho^{(k)}\mathbb{R}^{(k)}\right] \tag{D.11}$$

The likelihood maximisation procedure is then based upon iteratively diluting our estimate of ρ starting from the maximally mixed – unbiased – estimate $\rho^{(0)} = \mathbb{1}$. Note that the performance is sensitive to the choice of ε: one preceedure for setting this provided in reference [207] is to begin with $\varepsilon \to \infty$ (i.e. $\tilde{\mathbb{R}} \to \mathbb{R}$), and if the likelihood does not increase successively decrease ε to determine whether increases in likelihood are possible – if increases are possible, continue with the reduced value of ε, and once this appears to converge, vary ε to find the value for which any increase is maximised. Alternatively, they propose to randomly generate ε and keep attempting the iteration with different values until an increase in likelihood is seen, and move to the next step and repeat.

The authors note [207] that for most cases the $\varepsilon \to \infty$ converged and the likelihood increased monotonically; nevertheless this may "overshoot" in low dimensional systems.

Whilst not *optimal* we adopt an approximate analogy to the first approach and simply slowly

reduce $\varepsilon \to 0$ as the algorithm evolves. In this case we observe convergence and monotonic increase in the likelihood; although adding the random perturbation after the likelihod plateaus would be straightforward and reduce the probability of becoming trapped by a non-maximum-likelihood density matrix.

D.1.1 Extensions to MaxLik

If all system observables can be measured precisely, quantum tomography enables one to completely reconstruct an initially unknown state. If we can only access part of the system's observables and we have "informationally incomplete" data, it can be shown that in all but special cases, maximum likelihood estimators consistent with the data, are not unique [209], and within this family of solutions, there is no fundamental principle that dictates which of these to choose.

We note that in the presence of uncertainty MaxLik may tend to produce the most single mode estimate of ρ compatible with the observed data. In our case we wish to demonstrate *single modedness*, so having an algorithm which tends to produce a single mode result (within the family of states which are maximal in likelihood *subject to the informationally incomplete* measurement results) may be misleading.

To address this, one may consider a number of extensions to the present method. One example – discussed by the same authors as the above method – which may be applied is to maximise the entropy, and produce the "worst case" estimate *given the data* [233]. Including a step to maximise the entropy at each stage can be achieved by an additional Lagrange multiplier $\lambda\left[S(\rho) - S_{\max}\right]$, where $S(\rho) = -\mathrm{Tr}\left[\rho \log(\rho)\right]$ is the von Neumann entropy of the state, and $S_{\max}$ is the maximal value it may take within the possible space. Implementation of this requires additional computational overheads, and increases the runtime of the algorithm and introduces another tuning parameter (like the dilution parameter) which must be chosen appropriately. As such, this approach was not implemented in the results of the main text but remains an area of future work.

There are a range of such modifications one can consider, such as "hedging bets", in which one attempts to combat artifacts on the reconstruction due to *as yet unobserved* outcomes [234].

Alternatively, one might wish to go down the route of quantum Bayesianism [235]. In this case, one reformulates the problem, and considers the quantum state as the *belief* the observer has about a system, from which predictions can be made, rather than a statement of *fact* about the underlying system. The observer then updates their belief of the state based upon the resulting outcomes of measurements. From Bayes rule

$$p(H|D) = \frac{p(D|H)\,p(H)}{p(D)}, \tag{D.12}$$

where:

$p(D\|H)$	is the probability of observing outcome data D given a hypothesis H;
$p(H)$	is the *prior* probability that H is true;
$p(H\|D)$	is the *posterior* probability that H is true, given that we observed D;
$p(D)$	$= \sum_H p(D\|H)p(H)$ is the marginal probability of data D.

In this way, observing some measurement outcomes D updates our knowledge of the system – i.e. $p(H) \to p(H|D)$. In the quantum case the hypothesis H now corresponds to the belief about the quantum state ρ, which is then given by the posterior probability distribution $p(\rho)$ on the space of density operators, based upon one's current best knowledge of the system:

$$\rho = \int \mathrm{d}\rho' \rho' p(\rho'). \tag{D.13}$$

APPENDIX E

MODE SELECTIVITY MEASUREMENTS & NUMERICS

To supplement the results in chapter 7 I present some of the intermediate steps undertaken along the way.

E.1 Memory Measurements

E.1.1 Extracting Memory Efficiency

As in part I of this thesis, the measurement of the memory efficiency is extracted from statistics acquired from the APD counts resulting in traces of the form seen earlier in 3.11; I show typical figures for the carved input signals in figures E.1 for the COMP0 and COMP1 states at zero delay with respect to the control pulse.

The respective counts for each configuration are evaluated from integrating under the respective curves within the shaded regons.

Note that in the 100 ns after the initial signal pulse we observe a series of "echoes".[1] These are most likely effects from ringing of the amplifiers driving the FMZI, since measures were taken to eliminate any optical back reflections (replacing with APC fibres, and angling components from normal incidence). Similarly the relative magnitude of the echoes were unaffected by attenuating the signal incident on the APDs.

Furthermore, in the zoomed in plot of the retrieval time-bin (lower right figures of E.1) the same background structure is seen for "Memory on" and "Signal only", which is not present in the "Control only" trace.

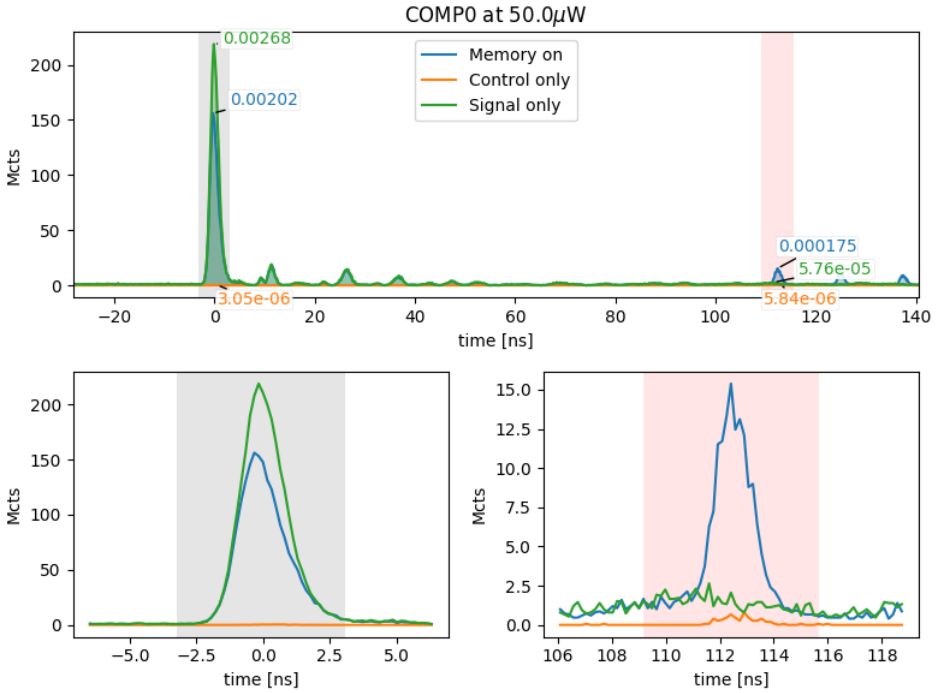
The evaluation of the memory efficiency is the same calculation as given in equation 3.141.

E.1.2 Time Scan Measurements

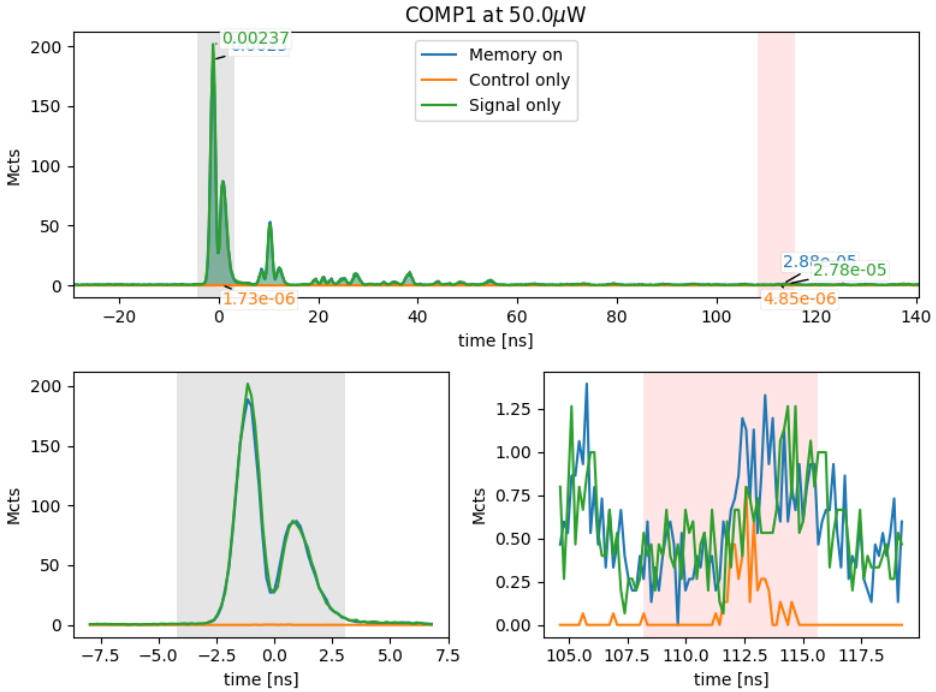
The results from the time–scanned 2 dimensional MUB states recorded at low incident control field power (as in figure 7.7) after rescaling, are replotted here in figures E.2 superimposed with the *time reversed*[2] probe state amplitudes and phase (inferred from electrode measurement) to provide a reference of the time scales.

We note a number of seemingly *repeated* features which not captured by the model used in the main text. One immediately noticeable feature is the peak near zero delay observed in each of the

[1] Displaying (i) periodic revivals, and (ii) exhibiting some of the features of the principle signal.
[2] Such that they follow the same trajectory as the efficiency as a function of delay.



(a) COMP0 input ∼ matched with control



(b) COMP1 input ∼ orthogonal to control

Figure E.1: APD counts over 10 s, for memory measurements with carved input signals. These were acquired at the "low" control field power in the main text; these are acquired with $\langle N_{\text{in}} \rangle \sim 9$, and attenuated after the memory.

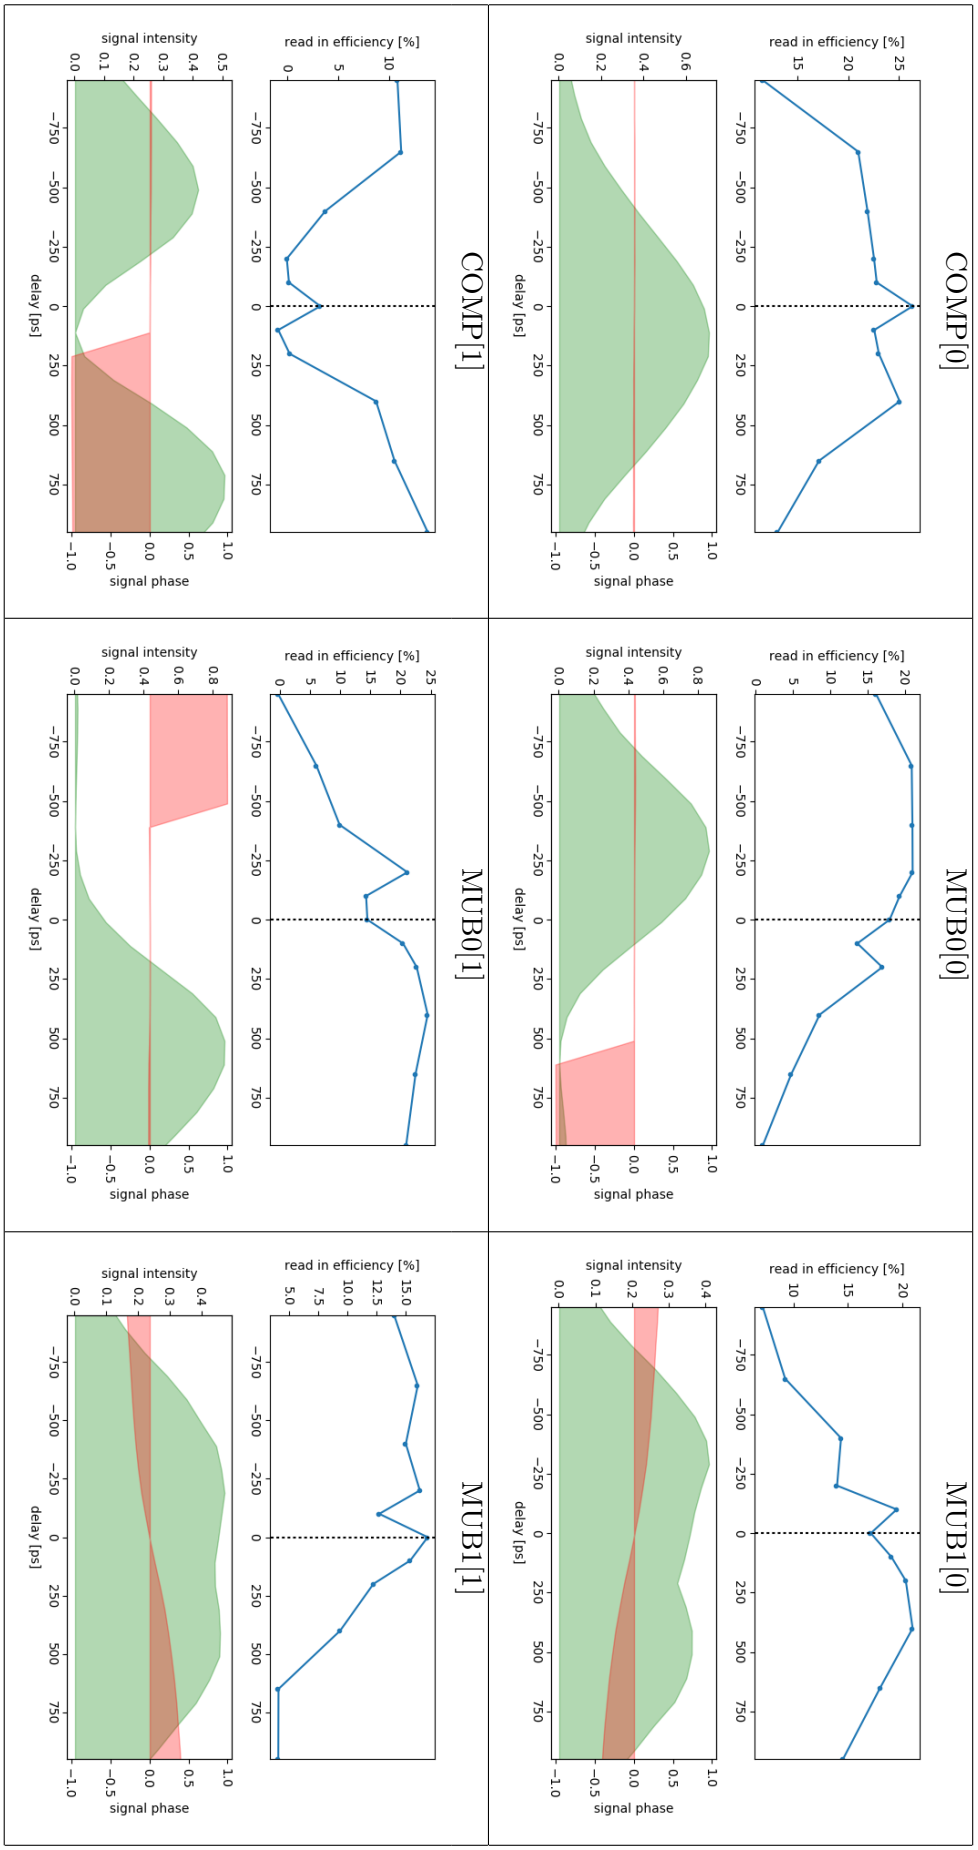


Figure E.2: Oddities in data: artifact or not?

traces, which whilst in consistently appears in this data set, it was not seen at the higher power, so is likely not a physical effect.

Secondly, whilst the magnitudes are different, there is an apparent functional similarity of the traces MUB01↔MUB10, and MUB00↔MUB11. Whilst intuitively the MUB0 efficiencies have the form as one might expect, the origin of this asymmetry in the MUB1 data are not obvious. We note that whilst the *inferred* preparation fidelity of the two MUB1 states was high ($\sim 96\%$) the optical intensities of these two states showed the worst overlap with the target states showing a significant mismatch in height.

One possibility investigated was whether the profile of the phase contributed to an increased efficiency of the MUB10 with respect to the MUB11 (since the "error function" phase profile is similar to that which is needed to combat the AC Stark shift; however this could not be reproduced in the (simplified) model. This difference was even greater in the higher power data and is not due to the difference in magnitude of the input signals, since this cancels out in the memory calculation.

E.2 Quadrature Methods

For the numerical evaluation of the integrals, a brief investigation of different sampling methods was performed.

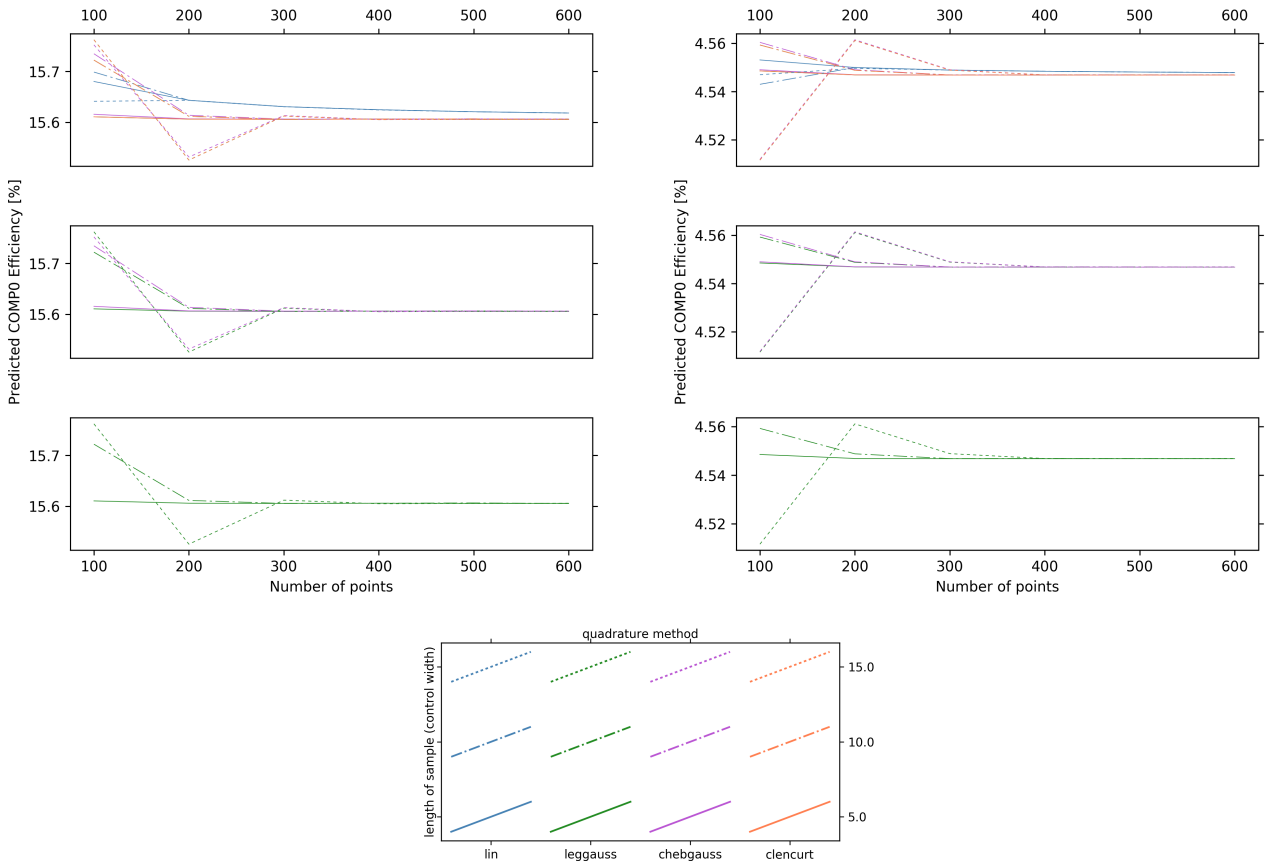


Figure E.3: (Top left/right) Convergence of memory efficiency with different quadrature methods as the number of sample points are increased (left/right are conducted at different control field powers) . (Bottom): different methods compared above: "lin" lenear sampling. "leggauss" Legendre-Gauss Quadrature, "chebgauss" Chebyshev-gauss, both evaluated using the numpy (numerical python) library "clencurt" Clenshaw Curtis, evalatuated using a python translation from scripts in [236]

The convergence of the simulated memory efficiency for the measured COMP0 state at zero delay, with a width matched control pulse, was compared for different quadrature methods as a function of the length of time window for different numbers of abcissas. We compare linear sampling with the Clenshaw–Curtis quadrature, and with Chebyshev– and Legendre– Gaussian quadratures [236]. With Gaussian quadratures, these are in general more sparsely sampled towards the centre of the domain, with more points towards the edges which also has the consequence of suppressing edge effects. All methods are in good agreement, with the Gaussian methods in general converging faster, and the linear being slowest.

After some investigation we chose to operate with a Legendre–Gauss quadrature with 512 points with time limits $\pm 10\times$ the estimated control width. The extended time window and number of points were chosen to accommodate for time–shifted probe states, broader than the control – since the tomography numerics were "better behaved" when the probe states returned to zero at the edges of the time window and satisfy pseudo–periodic boundary conditions, otherwise edge effects became apparent.

I also considered including a "super-Gaussian" filter function to damp waveforms to zero at the edges of the domain, however with the shorter window length this introduced other artifacts due to distorting the probe shapes. With the longer domain length $\sim 10\times$ the control width, the effects of the filter function depended on the basis of the reconstruction used, discussed in the next section.

E.3 Reconstruction Basis

I investigated implementing the tomography algorithm in the time domain – representing the interaction matrix directly in "physical units" – which necessarily involved the manipulation of large matrices, and was numerically intensive. Alternatively to reduce the dimensionality of the problem, I performed a coordinate transformation decomposing onto an approximately matched Hermite-Gaussian basis set (choosing 24 states, by trial and error). An example of this decomposition is shown in figure E.4 for a decomposition of one of the experimentally recorded probe states projected onto the HG basis (top right: amplitude, lower right: phase)for the MUB00 state, at zero delay.

Choosing a dimension of 24 enabled us to capture the timeshifted probe states without loosing information (as measured by the sum of the vector coefficients $\geq 99\%$ – shown in the legend of the figure E.4). This also naturally suppressed edge effects, and including the super-Gaussian filter to the probe states made no noticeable difference in the reconstruction.

The two choices of bases gave good agreement in the reconstructed states, and so in the main text we present the results from the decomposed method based on the significantly reduced computational overheads.

E.4 Tomography Algorithm

I monitor the evolution of the tomography algorithm by evaluating a metric of the Likelihood function, and the trace distance (i) between the ρ_k at subsequent iterations k of the algorithm and (ii) between the current ρ_k and the initial guess.

The typical trajectory observed is plotted in figure E.5, and in all cases, the algorithm exhibited reliable convergence.

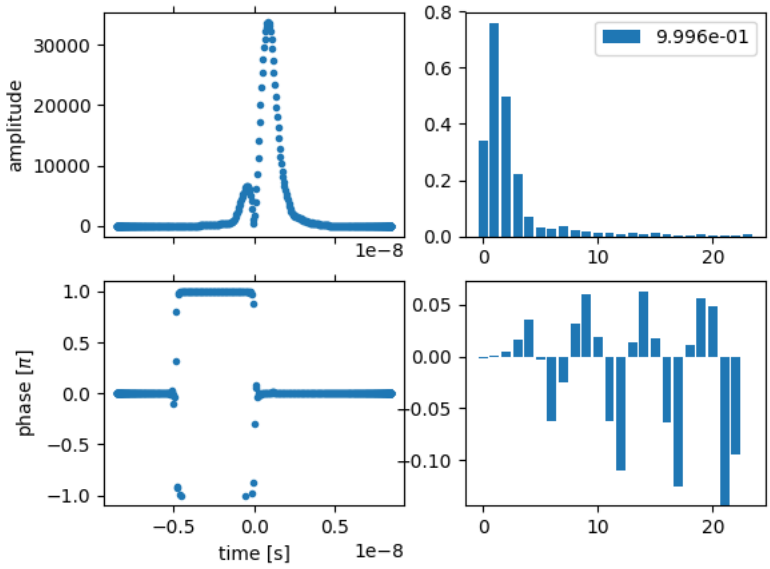


Figure E.4: Example of decomposition of probe state vector onto HG basis. (left) time domain, (right) HG basis, (top) amplitude, (bottom) phase. The legend gives the sum of the decomposition coefficients, $\sum |v_i|^2$, and indicates the "faithfulness" of the decomposition (i.e. a value 1 metric would mean no information is lost).

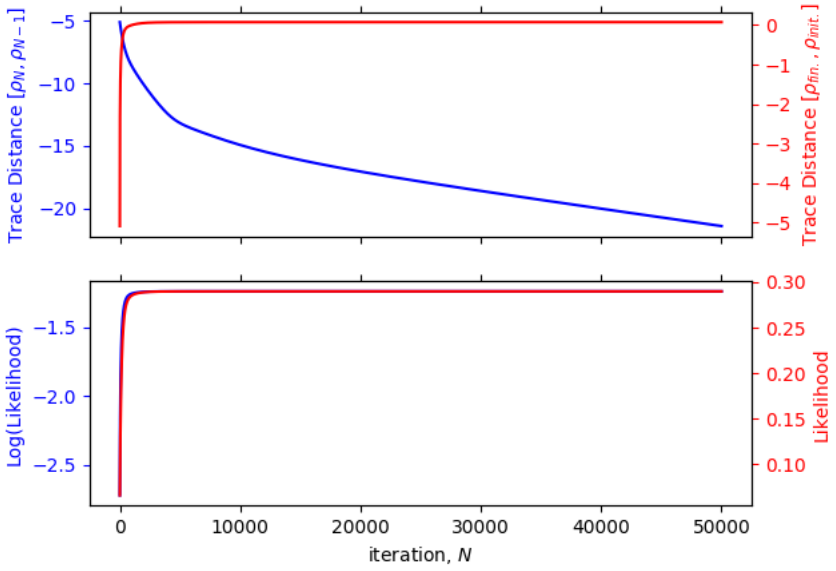


Figure E.5: Typical trajectory of the tomography algorithm convergence. The absolute value of the "likelihood" is not physically significant since, as evaluated within the execution the overall gain (hence resulting efficiency) is not known by the algorithm.

The convergence rate was sped up when the variable dilution parameter was included, without noteable difference in the final reconstructed state.

It is likely that there are grounds for further work in optimising the dilution parameters (and MaxEnt multiplier, if applicable).

thanks for reading

...pint, anyone?