

Storage and Manipulation of Optical Information Using Gradient Echo Memory in Warm Vapours and Cold Ensembles

Benjamin M. Sparkes

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

July 2013



This thesis is dedicated to my beautiful Mel,
who makes the world a better place simply by being in it.

Declaration

This thesis is an account of research undertaken between February 2009 and February 2013 at The Department of Quantum Science, Research School of Physics & Engineering, The Australian National University (ANU), Canberra, Australia. The author was a visitor at Laboratoire Kastler Brossel (LKB), Université Pierre et Marie Curie, Paris, France, from October 2011 to December 2011, and the work performed at the LKB is reported in Chapter 11.

The majority of the research presented here was supervised by Prof. Ping Koy Lam (ANU), Dr. Ben Buchler (ANU), and Dr. Thomas Symul (ANU). Part of the work recorded in Chapter 11 was supervised by Dr. Julien Laurat (LKB) and Prof. Elisabeth Giacobino (LKB). Except where acknowledged in the customary manner, the material presented in this thesis is, to the best of my knowledge, original and has not been submitted in whole or part for a degree in any university.

Benjamin M. Sparkes
July, 2013

Acknowledgements

“So Long, and Thanks for all the [Physics].” - Book Four in *The Hitchhiker’s Guide to the Galaxy* Trilogy, Douglas Adams, 1984, wilfully mis-quoted.

The work presented in this thesis would not have been possible without the help and support of many, many people.

I would like to start by thanking Dr. Mahdi Hosseini for all his help and advice over the past four years. Having someone as talented and knowledgeable as Mahdi around to start me on the right path was a large part of the reason why I was able to achieve the results presented here; while having his company in the lab also made the many hours spent there a much more pleasurable experience.

I must also thank my supervisors, Dr. Ben Buchler and Prof. Ping Koy Lam, for their assistance and guidance. Their innate understanding of both the experimental and theoretical components of quantum-atom optics is staggering, and I feel privileged to have caught even a sliver of this. I will always appreciate the freedom they allowed me to design and construct the cold atom GEM experiment.

A supervisor in all but name, I would like to express my sincerest gratitude to Dr. Nick Robins for all he has done for me during my time at ANU. It was working with Nick that gave me my first taste of quantum optics in 2006 and I haven’t looked back. I am still in awe at his generosity, both in terms of his equipment and his time: without his many cupboard-cows there is no way the cold atom GEM experiment could have been built in such a short time and his cold atom expertise was pivotal in achieving the results we did. In the last few months of my lab work, Nick’s endless optimism and amazingly positive personality not only kept me sane but also made the nights in the lab much less lonely.

A big thank you to all of the above who have given up their time to proof-read this thesis, especially Mahdi, which has made it a much more concise, comprehensible and complete work.

For my time in Paris I would like to thank Prof. Elisabeth Giacobino and Dr. Julien Laurat for giving me the opportunity to work on an exciting project in an amazing city. I learnt a lot from my visit and will always look back on my time at the LKB with fond memories. These were contributed to by Lambert Giner, Lucile Veissier, Adrien Nicolas and Sasha Sheremet, who were so welcoming and patient as I stumbled around both the lab and my French. *Merci bien!*

I have had the privilege of working with many exceptional people over the years. A big thank you to Dr. Julien Bernu for his help constructing the cold atom GEM experiment, his thoroughness and attention to detail helped make this project turn out as well as it did (not to mention his hard-core fibre aligning!). Thanks also to Jiao Geng for providing a dose of sunshine to the lab during the last few months of my PhD, and to Dr. Quentin Glorieux for his help with the final surge. I had the pleasure of working with Geoff Campbell, Daniel Higginbottom, Cameron Cairns, Dr. Gabriel Hétet and Dr. Olivier Pinel on the warm atom GEM experiment over the years, and would like to thank them for making the lab such a fun place to be. A special thank you to Helen Chrzanowski for

letting me experiment on her experiment, and for many interesting conversations. Thanks also to David Parrain for his assistance while working on the digital locking system and to Dr. Thomas Symul - without whose vision, and amazing knowledge of all things LabVIEW, the digital locking code would not have been created.

Other people whose help during my PhD was greatly appreciated include Paul Altin for his help on the cold atom experiment, Simon Bell for the variable-pitch solenoid design code, Daniel Döring for enlightened discussion on the ac Stark shift and his Mathematica code, and Pete Uhe for the initial testing of the multi-element coils. Thanks to all of you.

I would like to express my appreciation for the invaluable assistance I have received from various parts of the ANU over the years. Thanks to Shane Grieves, Neil Hinchey and James Dickson for their guidance on all things electronic and for building various critical components including LabVIEW break-out boxes, clock racks and the current drivers used for the spectral manipulation experiments. Thanks also to Oliver Thearle and Paul Redman for constructing the ADC racks to accompany the digital locking software, and to Colin Dedman for helping to design the fast switching GEM coils. There are too many experimental components made in-house that were vital to this thesis to ever list, but for their masterful creation I would like to thank Paul McNamara, Paul Tant and Neil Devlin from the mechanical workshop. I would also like to thank Gaye Carney, Laura Walmsley, Kerrie Cook and Amanda White for their help with all administration matters. A huge additional thanks to Amanda for the disproportionate amount of ordering she has placed for me, the fun conversations, and the occasional cookie.

While not contributing directly to the work presented in this thesis, there are many people without whom the last four years would have been far less enjoyable. Thanks to Phil and Richard for allowing me to disrupt their work on an almost-daily basis, I really enjoyed our conversations and the many, many Simpson quotes. My time in Paris was greatly enhanced by the presence of Mathieu and Fabbio, among others. Rock and roll guys! Back in Canberra I thank the rest of the quantum optics group for the fun times, and terrible jokes, we've had together.

To Mr. Owen and Rik, I thank you for the oases you provided away from the constant stream of physics, and your infinite patience when I invariably failed to practise. A huge thank you to Joe and Chrissie for all the lunches, TV nights (the fruits of which can be found at the start of almost every chapter), Brumbies matches and amazing dinners.

To Karen, Doug and Sarah, thank you so much for accepting me into your family, it is a great privilege and one that I will forever cherish. To my own family, Mum, Dad, Sophie and Jeremy, I will never be able to repay the belief you have always had in me or the unconditional support you have always offered. I count myself very lucky to have all of you in my life and, though "*thank you*" seems far too inadequate a phrase after all that you have done for me, I thank you from the bottom of my heart.

Lastly, and by far the most beautifully, to Mel - for always being there, for listening as I described my days in the lab in arduous detail (see Chapters 1-13), for understanding, for being happy for me when things went well, and for commiserating when things invariably didn't - thank you. But mostly thank you for being the amazing person that you are. I cannot wait to spend the rest of my life with you!

Abstract

Quantum memories for light lie at the heart of long-distance provably-secure communication [1], while containing the potential to help break current encryption methods [2], and allow better measurement of quantities than ever before [3]. Demand for a functioning quantum memory is therefore at a premium. Unfortunately, the same properties of light that make it such an effective carrier of quantum information make it difficult to store. Furthermore, by the laws of quantum mechanics, storage must be achieved without measurement to preserve the quantum state.

A quantum memory needs to have an efficiency approaching unity without adding noise to the state, and storage times from milliseconds to seconds. Ideally it would also have a high bandwidth and be able to store many pieces of information simultaneously. Many different techniques are currently being developed and much experimental progress has been made over the past few years, with: efficiencies approaching 90% [4]; storage times of over seconds [5]; bandwidths of gigahertz [6, 7]; and over 1000 pieces of information stored at one time [8]. These results were, however, achieved using different memory schemes in different storage media. The challenge now is to reproduce these results with one memory.

This thesis focuses on extending the gradient echo memory (GEM) scheme, which shows great promise due to the high efficiencies achieved (87%) [4]. GEM has also been used to demonstrate temporal compression and stretching of pulses, as well as a capacity to arbitrarily resequence stored information [9] and the interference of initially time-separated pulses [10].

Firstly, we demonstrate the noiseless nature of GEM storage in a warm vapour cell to prove that the output from the memory is the best-possible copy of the input allowed by quantum mechanics. We show GEM's ability to coherently and precisely spectrally-manipulate stored information by having fine control over the memory's frequency gradient, with potential applications for dynamic conditioning of information inside quantum networks [11]. We demonstrate cross-phase modulation of a stored light pulse with an additional optical field, a process with applications in quantum computing [12]. We also carry out storage of different spatial modes and arbitrary images, demonstrating the potential for orders of magnitude improvement in storage capacity.

We then switch from warm vapour cells to cold atomic ensembles to improve the storage time of GEM, seeing a maximum coherence time of 350 μs (seven times that of the warm vapour system) and achieving efficiencies of up to 80%, on a par with the highest efficiency achieved with a cold atomic ensemble [13]. In the process we developed an ultra-dense cold atomic cloud with potential applications in a range of quantum optics experiments. Cold atoms, and the small volumes they occupy, also allowed us to develop an alternative to using magnetic field gradients for our alkali-atom memories in the form of a light-field gradient. This holds promise for extremely fast gradient switching and fine control over the gradient.

We also present a digital locking code with application in a range of quantum optics experiments.

Contents

Declaration	iii
Abstract	vii
I Introduction to Quantum Memories	3
1 Introduction: Quantum Memories for Fun and Profit	5
1.1 Introduction to the Introduction	5
1.2 Quantum “Conundrums” and Resources	6
1.2.1 Quanta	6
1.2.2 Uncertainty, Superposition and Measurement	6
1.2.3 Phase and Coherence	7
1.2.4 Quantum Bits	7
1.2.5 Entanglement	8
1.2.6 Squeezing	8
1.3 Quantum Information Processing	9
1.3.1 Quantum Key Distribution, Quantum Repeaters, and Quantum Networks	9
1.3.2 Quantum Computing	12
1.3.3 Single-Photon Sources	15
1.3.4 Quantum Metrology	15
1.4 Quantum Memory Metrics	16
1.4.1 Efficiency	16
1.4.2 Fidelity	16
1.4.3 Storage Time	17
1.4.4 Bandwidth	17
1.4.5 Delay-Bandwidth Product	18
1.4.6 Multi-Mode Capacity	18
1.4.7 Wavelength	18
1.4.8 Storage Media	18
1.5 A Trip to the Quantum Memory Zoo	19
1.5.1 Delay Lines and Cavities	19
1.5.2 Electromagnetically-Induced Transparency	20
1.5.3 Raman Memory	21
1.5.4 The Duan-Lukin-Cirac-Zoller Protocol	22
1.5.5 Photon Echo Techniques	23
1.5.6 Other Species	30
1.6 Conclusions and Thesis Outline	31

2	Quantum Theory of Atom-Light Interactions	35
2.1	Introduction	35
2.2	Quantum Formalism	35
2.2.1	Quantum States and Operators	35
2.2.2	Quantum Measurements and Uncertainty	36
2.2.3	Pictures and Time-Evolution	37
2.3	Quantum Theory of Light	38
2.3.1	Number States, Creation and Annihilation Operators	38
2.3.2	Indistinguishability, Entangled States and Teleportation	38
2.3.3	Coherent and Squeezed States	41
2.3.4	The Quantised Multi-Mode Electric Field	42
2.3.5	Detection	44
2.4	Quantising Atoms	45
2.5	The Interaction of Atoms and Light	46
2.5.1	The Interaction Hamiltonian	46
2.5.2	The Dressed State Picture	48
2.5.3	Time-Evolution	48
2.5.4	Optical Bloch Equations	50
2.6	Conclusions	53
3	Quantum Memory Theory	55
3.1	A Memorable Introduction	55
3.2	Quantum Memory Metrics Take II	55
3.2.1	Efficiency, Fidelity and T-V Diagrams	55
3.2.2	Bandwidths and Coherence Times	57
3.3	Electromagnetically-Induced Transparency Theory	58
3.3.1	Slowing Light with EIT	58
3.3.2	EIT Polariton and Storing Light	59
3.4	Gradient Echo Memory Theory	60
3.4.1	Echo Memories	60
3.4.2	Two-Level Gradient Echo Memory	61
3.4.3	GEM Polariton	63
3.4.4	Stepping Up a Level - Three-Level GEM	65
3.5	Conclusions	67
4	A Digital Locking System for use on Quantum Optics Experiments	69
4.1	Introduction	69
4.2	Control Theory	71
4.3	The Digital Control System	75
4.3.1	The Hardware	75
4.3.2	The Software	77
4.3.3	The Controller	79
4.4	Locking Analysis and Optimisation	80
4.4.1	Measuring System Responses	80
4.4.2	Lock Optimisation	84
4.4.3	Long-Term Stability and Comparison with Analog PI Controller	85
4.4.4	Other Optimisations	85
4.5	Quantum Measurements	90

4.6	Conclusions	92
II	Warm Vapour Experiments	93
5	Introduction to Light Storage with Warm Atomic Ensembles	95
5.1	Introduction - Why Warm Atoms are Hot Stuff	95
5.2	Literature Review of Warm Atom Light Storage Experiments: 2001-Present	96
5.3	Not So Hot - Decoherence Effects in Warm Vapours	97
5.3.1	Doppler Broadening	98
5.3.2	Diffusion	99
5.3.3	Buffer Gas, Collisional Broadening and Anti-Relaxation Coating . .	100
5.3.4	Power Broadening and Spontaneous Raman Scattering	101
5.3.5	Four-Wave Mixing Effects	102
5.3.6	Magnetic Field Effects	103
5.4	Conclusions	104
6	Digitising GEM	107
6.1	Introduction	107
6.2	Λ -GEM Experimental Set-Up	107
6.3	Digital Acquisition for Noise Measurements	110
6.4	FPGA Control System	114
6.5	Digital Phase-Insensitive Heterodyne Demodulation	115
6.6	Conclusions	116
7	Precision Spectral Manipulation Using GEM	117
7.1	Introduction	117
7.2	The Multi-Element Coil	119
7.3	Spectral Manipulation Experiments	122
7.3.1	Centre Frequency Offset Characterisation	122
7.3.2	Input/Output Bandwidth Characterisation	125
7.3.3	Spectral Filtering	128
7.3.4	Pulse Interference	130
7.4	Cross-Phase Modulation	133
7.5	Discussion	135
7.6	Conclusions	138
8	Spatially Multi-Mode Storage Using GEM	139
8.1	Introduction	139
8.2	Spatial Mode Detection Set-Up	140
8.3	Four-Wave Mixing Pulse	142
8.4	Spatial Mode Storage Results	143
8.5	Selective Storage	145
8.6	Image Storage	145
8.7	Discussion	145
8.8	Conclusions	147
III	Cold Atom Experiments	149

9	Introduction to Light Storage with Cold Atomic Ensembles	151
9.1	Introduction - Why Cool Kids Use Cold Atoms	151
9.2	Theory of Atom Trapping and Cooling	152
9.2.1	Magneto-Optical Traps	152
9.2.2	Further Cooling	155
9.2.3	Dipole Traps and Optical Lattices	156
9.3	Decoherence Effects in Cold Atomic Ensembles	159
9.3.1	Collisions	159
9.3.2	Diffusion and Time of Flight	159
9.3.3	Scattering Rate	160
9.3.4	Heating Rate	162
9.4	Literature Review of Cold Atom Light Storage Experiments	162
9.4.1	A History of Cold Atom Light Storage Experiments 1999-Present . .	162
9.4.2	Experimental Improvements	164
9.5	Conclusions	166
10	ac Stark Gradient Proposal	167
10.1	Introduction - An Enlightened Idea	167
10.2	AC Stark Shift Proposal	169
10.2.1	AC Stark Shift Theory	169
10.2.2	Atom Trapping	172
10.2.3	Gradient Creation	173
10.2.4	Switching Protocols	175
10.3	Limiting Factors	178
10.3.1	Time Scales	178
10.3.2	Efficiency	179
10.4	Experimental ac Stark Gradient Generation	180
10.5	Conclusions	182
11	Experiments Towards Cold Atom GEM	183
11.1	Introduction	183
11.2	Experiment 1 - Three-Tier Rubidium MOT and Dipole Trap	183
11.2.1	Memory Set-Up	184
11.2.2	Atom Trapping Set-Up	185
11.2.3	Absorption Imaging	188
11.2.4	Interface and Detection Set-up 1a - Opposite Polarisation, $F = 2$.	191
11.2.5	Interface and Detection Set-up 1b - Opposite Polarisation, $F = 1$.	191
11.2.6	Interface and Detection Set-up 1c - Same Polarisation, $F = 1$. . .	195
11.2.7	Interface and Detection Set-up 1d - Dipole Trap, $F = 1$	199
11.2.8	Experiment 1 Discussion	201
11.3	Experiment 2 - Caesium MOT	202
11.3.1	Experiment 2 Set-Up	202
11.3.2	Diagnostic Tools	204
11.3.3	EIT Experiments	207
11.3.4	Experiment 2 Discussion	210
11.4	Conclusions	211

12 Cold Atom Gradient Echo Memory	213
12.1 Introduction	213
12.2 Common Experimental Set-Up	213
12.2.1 Table 1	213
12.2.2 Table 2	214
12.3 Configuration 1 - MOT Fields at 45° to GEM Fields	215
12.3.1 MOT 1 Set-Up, Optimisation and Characterisation	215
12.3.2 GEM Set-Up and Results with MOT 1	218
12.4 Configuration 2 - MOT Fields Parallel and Perpendicular to GEM Fields	221
12.4.1 MOT 2 Set-Up, Optimisation and Characterisation	221
12.4.2 GEM Set-Up and Results with MOT 2	223
12.5 Configuration 3 - MOT Fields at 45° to GEM Fields, Again	225
12.5.1 MOT 3 Set-Up, Optimisation and Characterisation	225
12.5.2 GEM Set-Up and Results with MOT 3	228
12.6 Discusion	231
12.7 Conclusion	232
13 Conclusions and Future Outlook	237
A Alkali Structure	241
A.1 Comments on Level Structures	241
A.2 Rubidium	241
A.3 Caesium	244
B PDH Lock 5001 Instruction Manual - Rev. 0.1	245
B.1 Introduction	246
B.2 Set-up	246
B.2.1 Hardware	246
B.2.2 Software	248
B.3 First Time	248
B.4 Locking the System	249
B.5 Lock Optimisation	252
B.5.1 Control Theory	252
B.5.2 Lock Optimisation Protocol	254
B.6 Sequential Locking	256
B.7 Menu Overview	258
B.7.1 Lock #	258
B.7.2 Clocks and Timing and Misc	261
B.7.3 Sequential Locking	263
B.7.4 Scope	263
B.7.5 Transfer Functions	264
B.7.6 Controller Measures	266
B.7.7 Saving Data	267
B.7.8 Re-ordered IP/OPs and Errors	268
B.7.9 PC_Sound_Gen.vi	269
B.8 Known Problems	269
B.9 Acknowledgements	270

C	Cavities and Pound-Drever-Hall Locking	271
C.1	Cavities	271
C.2	Pound-Drever-Hall Locking	273
D	Mathematica Codes	275
D.1	Code to Calculate Multi-Element Coil Currents	275
D.2	Heterodyne Data Extraction Code	278
D.3	ac Stark Shift Code	284
	Bibliography	289

List of Figures

1.1	Classical Vs Quantum Memory	7
1.2	Qubit	8
1.3	Squeezed State	9
1.4	Quantum Repeater Protocol	11
1.5	Quantum Computing	13
1.6	Single-Photon Gun	15
1.7	Fidelity	17
1.8	Electromagnetically-Induced Transparency	20
1.9	Raman Memory	21
1.10	DLCZ Scheme	23
1.11	Two-Level Photon Echo Scheme	24
1.12	Three-Pulse Photon Echo Scheme	25
1.13	AFC Scheme	26
1.14	CRIB Scheme	27
1.15	The GEM Scheme	29
1.16	Thesis Outline	33
2.1	Indistinguishability	39
2.2	Coherent, Squeezed and Thermal States	40
2.3	Transverse Electromagnetic Modes	43
2.4	Heterodyne Detection Set-Up	44
2.5	The Three-Level Atom	46
2.6	Dressed States	47
2.7	Bloch Sphere	50
3.1	Loss from Quantum Systems	55
3.2	GEM Bandwidth	62
3.3	GEM Polaritons	64
3.4	Two-Level, Three-Level Equivalence	66
4.1	Feedback Control System	72
4.2	Controller Properties	73
4.3	Digital Locking System Set-Up	76
4.4	Digital Locking Racks	77
4.5	Sequential Locking	78
4.6	PII Controller	79
4.7	Approximating G and Simulating H	81
4.8	Lock Optimisation Protocol	83
4.9	Analog and Digital Controller Comparison	86
4.10	Low-Pass Filtering	87
4.11	Digital Noise	89

4.12	Schrödinger Cat Experiment	91
5.1	Warm Vapour Cell	95
5.2	High Efficiency and Multi-Pulse Λ -GEM Storage	97
5.3	Diffusion and Buffer Gas	100
5.4	Four-Wave Mixing Double- Λ Scheme	103
6.1	Λ -GEM Set-Up	108
6.2	Zeeman States	109
6.3	Λ -GEM Noise Measurements 1	111
6.4	Λ -GEM Noise Measurements 2	112
6.5	FPGA Control and Monitoring	113
6.6	Digital Phase-Insensitive Heterodyne Demodulation	115
7.1	Multi-Element Coil Set-Up	120
7.2	LabVIEW MEC Gradient Code	121
7.3	Basic Λ -GEM Set-Up	122
7.4	Centre Frequency Offset Characterisation	123
7.5	Extra Centre Frequency Offset Characterisation	124
7.6	Input/Output Bandwidth Characterisation	126
7.7	Extra Input/Output Bandwidth Characterisation	127
7.8	Spectral Filtering	129
7.9	Interference with Pulses of Different Frequencies	131
7.10	Interference with pulses of the same frequency	132
7.11	Cross-Phase Modulation Set-Up	134
7.12	Cross-Phase Modulation Results	136
7.13	Memory Network	137
8.1	Spatial Mode Detection Set-Up	141
8.2	Four-Wave Mixing Pulse	142
8.3	Multi-Mode Storage in GEM	144
8.4	Selective Maxima Storage	145
8.5	Image Storage	146
9.1	Magneto-Optical Trap	153
9.2	Dipole Trap Modelling	157
9.3	Cold Atom Diffusion	160
9.4	Scattering Rate	161
10.1	Proposed ac Stark GEM Experiment	168
10.2	Rubidium-87 Level Structure	169
10.3	ac Stark Frequency Splitting	171
10.4	Scattering Rate for a One Megahertz ac Stark Gradient	173
10.5	ac Stark Switching Methods	177
10.6	ac Stark Cold Atom GEM Efficiency	178
10.7	ac Stark Gradient Generation	181
11.1	Cold Atom Memory Set-Up, Expt. 1	184
11.2	Cold Atom Trapping Set-Up, Expt. 1	186
11.3	Interface and Detection Set-up 1a,b	190

11.4	Scanning EIT in MOT on $F = 2$, Expt. 1a	192
11.5	Population Pumping from $F = 2 \rightarrow F = 1$ Level, Expt. 1b	193
11.6	EIT in MOT on $F = 1$, Expt. 1b	194
11.7	Interface and Detection Set-up 1c, and Optical Pumping Effects	196
11.8	Optimising Magnetic Field Compensation, Expt. 1c	197
11.9	Raman Absorption, Expt. 1c	198
11.10	Interface and Detection Set-Up 1d - Dipole Trap	200
11.11	Cold Atom Cs Memory Set-Up, Expt. 2	203
11.12	Optical Depth Measure, Expt. 2	205
11.13	Magnetic Field Compensation, Expt. 2	206
11.14	EIT Delay and Storage, Expt. 2	207
11.15	EIT to Autler-Townes Transition, Expt. 2	209
12.1	Cold GEM Set-Up 1	216
12.2	Imaging Frequency Calibration	217
12.3	MOT 1 Characterisation Images	218
12.4	MOT 1 GEM Results	220
12.5	Cold GEM Set-Up 2	221
12.6	MOT 2 Characterisation Images	222
12.7	GEM Coil Switching	223
12.8	MOT 2 GEM Results	224
12.9	Cold GEM 3 Set-Up	225
12.10	MOT 3 Characterisation Images	226
12.11	Peak OD Characterisation	227
12.12	Probe Diffraction	228
12.13	MOT 3 GEM Results	230
12.14	Full Cold Atom GEM Set-Up - Table 1	233
12.15	Full Cold Atom GEM Set-Up - Table 2.1	234
12.16	Full Cold Atom GEM Set-Up - Table 2.2	235
12.17	Full Cold Atom GEM Set-Up - Table 2.3	236
A.1	Full Rubidium-87 Level Structure	242
A.2	^{87}Rb Saturation Absorption Traces	243
A.3	Caesium-133 D2 Level Structure	244
B.1	Manual - Set-Up	247
B.2	Manual - First Time	249
B.3	Manual - Locking the System	250
B.4	Manual - Lock Optimisation	253
B.5	Manual - Lock Optimisation Protocol	255
B.6	Manual - Sequential Locking	257
C.1	Cavity Theory	272
C.2	PDH Error Signals	273

List of Tables

1.1	Quantum Memory Performance Summary	31
5.1	Decoherence Rates and Frequency Changes in Warm Vapours	104
7.1	Interference Echo Parameters	133
11.1	Main Run Times for Expt. 1	188
11.2	Beam Alignment Run for Expt. 1	189
11.3	Light Storage Runs for Expt.1a-d	192
11.4	Light Storage Run for Expt.1d (Dipole Trap)	199

Part I

Introduction to Quantum
Memories

Introduction: Quantum Memories for Fun and Profit

Chuck: “Hey, ... here are a couple of things you might need to know, or maybe you just forgot.” - “Chuck Versus the Ex” Chuck, NBC, 10 November 2008, TV series.

1.1 Introduction to the Introduction

Quantum information processing (QIP) encompasses the fields of quantum key distribution, quantum computing and quantum metrology. What draws all these fields together is their combination of quantum principles, such as superposition and entanglement, with information theory. I feel this is best articulated by a quote from N. Gisin et al. in their review of quantum key distribution in 2002 [14]:

“[O]ne could characterise quantum information processing as the science of turning quantum conundrums into potentially useful applications.”

For those potentially useful applications quantum information processing offers great promise, from providing provably-secure long-distance communication, while at the same time offering the potential for breaking our current, classical, security algorithms. However, to become experimentally operational, many aspects of QIP require a quantum memory - a device that can store quantum states without measuring them. Not only would a quantum memory have practical uses, but the study of light-matter interactions entailed in their development are of fundamental scientific interest.

It is impossible to achieve the functionality of a quantum memory by classical, measure-and-reproduce, means due to the “conundrum” of wave-function collapse. In the last decade or so much progress has been made towards the development of a quantum memory and in this chapter we will review the main methods used and results achieved in Sec. 1.5. But first we will discuss the main quantum “conundrums” in Sec. 1.2 and motivate the need for quantum memories by examining its uses in quantum information processing in Sec. 1.3. We will also discuss what makes a “good” quantum memory in Sec. 1.4 and finish the chapter with an overview of the best results achieved from quantum memories in Sec. 1.6, as well as outlining the structure of the rest of the thesis.

1.2 Quantum “Conundrums” and Resources

In this section we will briefly review some of the “odd” properties of quantum mechanics alluded to by Gisin et al. in the quote above, most of which are resources relevant for quantum information processing. For a more mathematical treatment of some of these resources see Ch.s 2 and 3.

1.2.1 Quanta

The name “*quantum mechanics*” derives from the fact that, unlike classical mechanics, everything in quantum mechanics is discretised into quanta. The most famous example of this is the photo-electric effect, where Einstein postulated that light was made up of individual particles, now termed photons, that carried a set amount of energy proportional to their frequency [15]. Bohr also showed that the internal energy levels of atoms are quantised, explaining their discretised absorption and emission spectra [16]. These are the main quantisation phenomena that we will make use of in this thesis.

1.2.2 Uncertainty, Superposition and Measurement

Another major difference between classical and quantum mechanics is that we can never know everything about an arbitrary quantum state. This stems from the Heisenberg uncertainty principle which states that certain pairs of system properties (most famously position and momentum) cannot be measured simultaneously past a certain uncertainty limit [17].

A quantum system is also probabilistic rather than deterministic. This means that a quantum system can be in multiple states at the same time. However, if we were to measure the system we would always find it in one state as the act of measuring a quantum system will cause its probability amplitude, or wavefunction, to collapse into one of the possible states. Therefore measurement will alter a quantum system. An example that will be of use later in this chapter is polarisation of light. Let us suppose that the polarisation of the light is at 45° to the horizontal (i.e., diagonal - D). Mathematically we can treat this as the light being in a superposition of horizontal (H) and vertical (V) polarisation states with equal probability, i.e., $|D\rangle = |H\rangle + |V\rangle$ (ignoring normalisation). If we then measure the polarisation of the light in this horizontal-vertical basis, we will find that 50% of the time it will be horizontal and 50% of the time it will be vertical. This is illustrated in Fig. 1.1(a)(i) and can be used to create a quantum random number generator [18].

The fact that measurement will perturb the quantum state has deep implications for quantum information processing. For instance, it led W. K. Wootters and W. H. Zurek to determine that single quanta cannot be cloned [19]. This, in turn, leads to the differentiation between classical and quantum memories. Classical memories use measurement and reproduction (i.e., cloning) to faithfully store quantum states. Quantum memories, on the other hand, must use anything but measurement to localise a quantum state for a controllable period of time. This difference is illustrated in Fig. 1.1.




Figure 1.1: Classical Vs Quantum Memory. (a) A classical memory with (i) measurement (M) and (ii) reproduction (R). This scheme will only work for a set of orthogonal states. If this is not the case, as here with the diagonal (D) photons being measured with a horizontal-polarisation filter (H Filter) - a superposition of horizontal (H) and vertical (V) photons as shown in the inset - then the superposition will collapse upon measurement, thus removing quantum information. (b) A quantum memory (QM) (i) stores information, without measurement, and (ii) retrieves the information when triggered (T).

1.2.3 Phase and Coherence

Not only can quantum systems be in a multitude of states at the same time, as discussed above, but there is also a phase between these states. This is why, when G. I. Taylor performed his famous single-photon double-slit experiment in 1909, he was able to see interference at the output [20]¹. If the phase between the two states (i.e., left- and right-slit states) was not well defined, then the interference pattern would not be visible. This is an illustration of the importance of the preservation of phase, otherwise known as coherence, in quantum systems. For instance, if the phase between the superposition of horizontal and vertical polarisation states in Fig. 1.1(b) was lost inside the memory, it could only perform as well as its classical equivalent. Another, rather dramatic, illustration of the link between quantum systems and coherence was the experiment by M. R. Andrews et al., who showed interference between two macroscopic ensembles of atoms in a Bose-Einstein condensate, a state of matter where all atoms occupy one quantum state [22].

1.2.4 Quantum Bits

Combining the quantisation, superpositions and phase properties led to the development of the quantum bit, or qubit. In classical computing a bit is a piece of information that can be either a 0 or a 1. In quantum mechanics, however, we can have a system which is partly in a 0 state and partly in a 1 state, with some relative phase between them. This gives an idea of how useful quantum mechanics can be for improving computation processes.

To illustrate all the information contained in a qubit we can use a unit sphere, called the Bloch sphere, where the height of the vector represents the relative probability of being

¹It is now known that this type of experiment could not in fact prove interference between individual quanta. The first experiment to successfully demonstrate this was by P. Grangier et al. in 1986 [21].

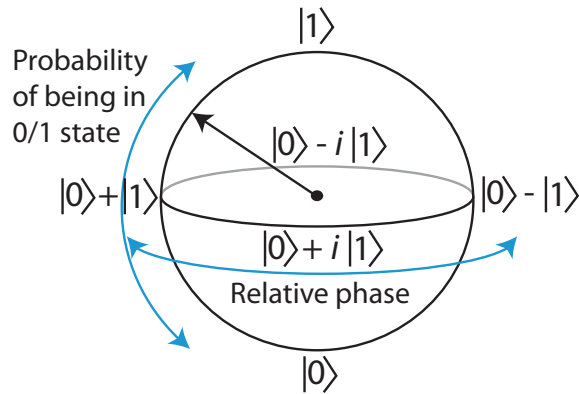


Figure 1.2: Qubit. A Bloch sphere representation, with the probability of the qubit being in state 0 or 1 given by how high the vector points and the relative phase between the two states given by the azimuthal angle (superposition states not normalised).

in state $|1\rangle$, as opposed to state $|0\rangle$. The azimuthal angle, meanwhile, tells us about the relative phase between the two. This is shown in Fig. 1.2.

Experimentally, qubits can be superpositions of polarisation, spatial-mode or time-bin states [23], or other states such as photon-number [24] or path [25]. Normally photons are used as qubits as they preserve their state for long times due to their lack of interaction with the environment. They also allow travel at the speed of light, as fast as any information is allowed to travel in our universe, and therefore allow for fast data transfer.

1.2.5 Entanglement

Entanglement is the quantum mechanical phenomenon whereby two quantum states are correlated in a non-local manner. This means that if, for instance, one of the two entangled particles is measured, it will also cause the *instantaneous* collapse of the other particle's wavefunction. Like most counter-intuitive quantum mechanical effects, early critics of quantum theory used entanglement to conclude that the theory was not complete [26]. Nearly 80 years later, however, entanglement is now a key component of quantum teleportation [27], quantum key distribution [28] and quantum computing protocols following experiments by A. Aspect et al. in 1981 [29].

Pairs of entangled photons are, these days, generally created via type-II spontaneous parametric down-conversion, first demonstrated by P. G. Kwiat et al. in 1995 [30]. This is where a strong pump laser is sent into a non-linear crystal. There is then a possibility that pump photons will split into two photons of roughly half the frequency, being entangled in momentum, energy and polarisation. This is a probabilistic process.

1.2.6 Squeezing

The Heisenberg uncertainty principle, discussed above, gives us a bound on how much information we can simultaneously know about certain pairs of properties of a quantum system. This is known as the standard quantum limit for pairs of properties such as energy and momentum, or amplitude and phase. This does not mean, however, that we can't




Figure 1.3: Squeezed State. Ball-on-stick representation of a squeezed state of light, with dashed blue circle showing the initial, symmetric, Heisenberg uncertainty limit of a coherent state (see Sec. 2.3.3), with an amplitude determined by the size of the stick. The solid green ellipse shows the phase-squeezed final state, with the arrows indicating the direction of squeezing.

make a measurement below the quantum limit, as long as the overall uncertainty remains the same (i.e., an increase in uncertainty in the other observable). This is known as squeezing (see Ref. [31] and other references in the same issue), as the circular Heisenberg limit is squashed into an ellipse, as illustrated in Fig. 1.3.

Squeezing, first demonstrated by R. E. Slusher et al. in 1985 [32] using four-wave mixing (see Sec. 1.5.6), is useful for precision metrology applications where only one of the two properties is of interest, for instance the phase in a gravitational-wave interferometer (see Ref. [33] and references therein). Some of the best squeezing results have been demonstrated using optical parametric processes, with up to 12.7 dB of noise suppression below the Heisenberg limit at high frequencies [34] and 10 dB down to 10 Hz [35].

1.3 Quantum Information Processing

We will now discuss real-world applications of these quantum phenomena in quantum information processing, focusing on the role played by quantum memories.

1.3.1 Quantum Key Distribution, Quantum Repeaters, and Quantum Networks

Quantum key distribution (QKD) is perhaps the first practical use of single-photon quantum phenomena. The original QKD proposal came from C. H. Bennett and G. Brassard in 1984, who suggested using two non-orthogonal bases for polarised photons (i.e., horizontal and vertical, and diagonal and anti-diagonal) to create an encryption key for a one-time pad [36]. It had already been shown, by C. Shannon in 1949, that if a key was completely random and the same length as the message to be encoded, then it could not be broken

[37]. However, without QKD there are practical issues with attempting to safely distribute such keys over large distances (i.e., trusting couriers etc.).

Bennett and Brassard's idea was to use quantum information encoded onto photons to create a key between two parties, Alice and Bob. For instance, if Alice wanted to send an encoded message to Bob, they would first create their secret key. To do this, Alice encodes a 0 or a 1 onto the polarisation of a photon in one of the bases, for instance the horizontal-vertical basis. She then sends this photon to Bob who proceeds to measure it in one of the two bases. If Alice and Bob use the same basis (determined through a classical communication channel after Bob's measurements) then they can say with *certainty* that the bit was either a 1 or 0 and include it in their secret key. However, if they measure in different bases then, due to the measurement Bob performed in the "wrong" bases, the photon he received will have collapsed with equal probability into a 0 or 1 (see Fig. 1.1(a)). They therefore discard these photons from their sifted key. Once they have enough bits measured in the same basis they then have their shared key that Alice can use to encode her information to send to Bob over a classical channel, and Bob can use the same key to decode the message.

The fact that measurement affects the quantum state is also important if any eavesdropper, Eve, decides to try and spy on Alice and Bob (as is her wont). To be able to glean information, Eve will have to measure the quantum state being sent from Alice to Bob. However, if she does so, there is no way she can make an identical copy of the state due to the no-cloning theorem [19]. Therefore, if Alice and Bob were to compare some bits of their sifted key and noticed large discrepancies between the 1s and 0s encoded and measured, they could deduce the presence of Eve. The proof of security for this scheme was proved by P. W. Shor and J. Preskill in 2000 [38].

An alternate QKD protocol was developed by A. K. Ekert in 1991 where, instead of Alice sending a pre-prepared photon to Bob, a polarisation-entangled photon source would be placed between them and one photon sent to Alice and one to Bob [28]. This is shown in Fig. 1.4(a). The beauty of this scheme is that, by testing the non-classicality of the correlations between the photon pairs, Alice and Bob can guarantee that Eve is not present.

The first experimental demonstration of QKD was over 32 cm of free-space in 1992 by C. H. Bennett et al. [39]. Since then distances up to 144 km in free-space [40, 41], and between 250-260 km in fibres [42, 43] have been achieved. Apart from single-photons, and entangled photon pairs, QKD has also been demonstrated with continuous variables (see review in Ref. [44]).

Extending the distance of QKD is problematic as quantum states are sensitive to loss and noise in the system. These will include dark counts and inefficiencies of the detectors but, in terms of scaling, the biggest issue is loss in the quantum channel. For instance, at its optimal wavelength of around 1500 μm , optical fibre has a loss of approximately 0.2 dB km⁻¹ [45]. This means that, for a total distance of 1000 km and an initial key creation rate of 1 GHz, it would take on average over 3000 years for a photon to reach the other end. Obviously this is far too slow for any practical purpose.

Meanwhile, free-space communication is limited by line-of-sight and the weather. To a certain extent these limitations can be overcome with combination of ground bases and satellites [46, 47]. However, apart from the prohibitive cost, if measurement were to occur on the satellites it would involve a third party, Charlie, who would be needed to create one key with Alice and one with Bob. This would provide Eve with a chance to again use nefarious classical means to obtain Alice's and Bob's secret message.




Figure 1.4: Quantum Repeater Protocol. (a) Ekert QKD protocol, where one each of a pair of entangled photons (1 and 2) are sent separately to Alice and Bob, who perform measurements on them in non-orthogonal bases. Key rates reduce exponentially with distance. (b) Simplified quantum repeater protocol - (i) the distance between Alice and Bob is divided by placing quantum repeater (QR) nodes containing quantum memories (QMs) between them. Multiple entangled-photon pairs 1,2 and 3, 4 and 5,6 are probabilistically generated and deterministically stored in QMs. (ii) Bell-state measurement (B) is performed at one node, and the result is fed forward classically (CFF) to the next node to teleport the quantum state of Photon 5 onto Photon 3. (iii) The process is repeated until the quantum state of Photon 5 has been teleported to Photon 1. This is then equivalent to (a), but with a polynomial decrease in key rate with distance.

To attempt to overcome the exponential loss of information and avoid any third party involvement, the idea of a quantum repeater was proposed by H.-J. Briegel et al. in 1998 [48]. Their idea was that, by dividing up the distance from Alice to Bob into many sub-lengths separated by nodes, the signal decay could be made polynomial with distance as opposed to exponential. The basics of the quantum repeater protocol are illustrated in Fig. 1.4(b). A quantum repeater is made up of three parts: an entangled (Bell-state) photon source, a detection set-up; and a quantum memory. In the first step of the quantum repeater protocol, a pair of entangled photons is separated and sent to different quantum repeater nodes and stored in the quantum memories located there. Once two photons from two different entangled pairs (i.e., Photons 4 and 5 from Fig. 1.4(b)) are stored in a node, they are removed from the memory and a Bell-state measurement is performed [49, 50]. This measurement gives information about the relationship between the two photons measured. This information is then sent through a classical channel from one quantum repeater to the next so that operations can be performed on the second photon from one of the pairs in the first repeater to make it indistinguishable to the photon from the other pair (i.e., to make Photon 3 indistinguishable from Photon 5). This is called quantum teleportation, first demonstrated by C. H. Bennett et al. in 1993 [27]. By this method entanglement can be swapped down the quantum repeater chain from Bob to Alice [51], with Photon 3 being teleported to Photon 1 so that we end up with Photons 1 and 6 being entangled. These photons are then retrieved from their quantum memories and sent to Alice and Bob, equivalent to the Ekert protocol.

Without quantum memories, this protocol would not scale polynomially due to the probabilistic nature of entangled-photon-pair generation. If this were the case, the protocol would only work if all photon sources emitted a pair of photons at exactly the same time, and even then we are assuming 100% efficient detectors and communication lines. Quantum memories mean that once a photon pair is produced it can be stored until required. And, if a detection fails, then not all entangled pairs need be reproduced. Many different quantum repeater protocols have been developed, with the most famous being the Duan-Lukin-Cirac-Zoller (DLCZ) protocol [1], which will be discussed in more detail in Sec. 1.5.4. Most protocols also involve some form of entanglement purification [52].

Apart from the goal of globally provably-secure communication via QKD, quantum memories will also allow for the development of quantum networks. Such a network will consist of quantum-repeater-style interconnected nodes, where quantum memories are used to store information, which is then processed and teleported along as necessary [53, 11]. These networks will have the potential to allow distributed quantum computing, as well as multi-party quantum communication that could be put to such uses as quantum secret sharing [54].

For more information on QKD see review articles Ref.s [14] and [55].

1.3.2 Quantum Computing

Quantum computing, i.e., computing using qubits instead of bits, could offer significant advantages over its classical counterpart. In particular, there are certain problems that could be solved much more efficiently (i.e., polynomial resources as opposed to exponential). These include L. K. Grover's algorithm for database searching [57, 58], P. W. Shor's algorithm for factoring large numbers into two primes [2], as well as the ability to simulate quantum dynamics [59]. The second of these three abilities has, perhaps, the most

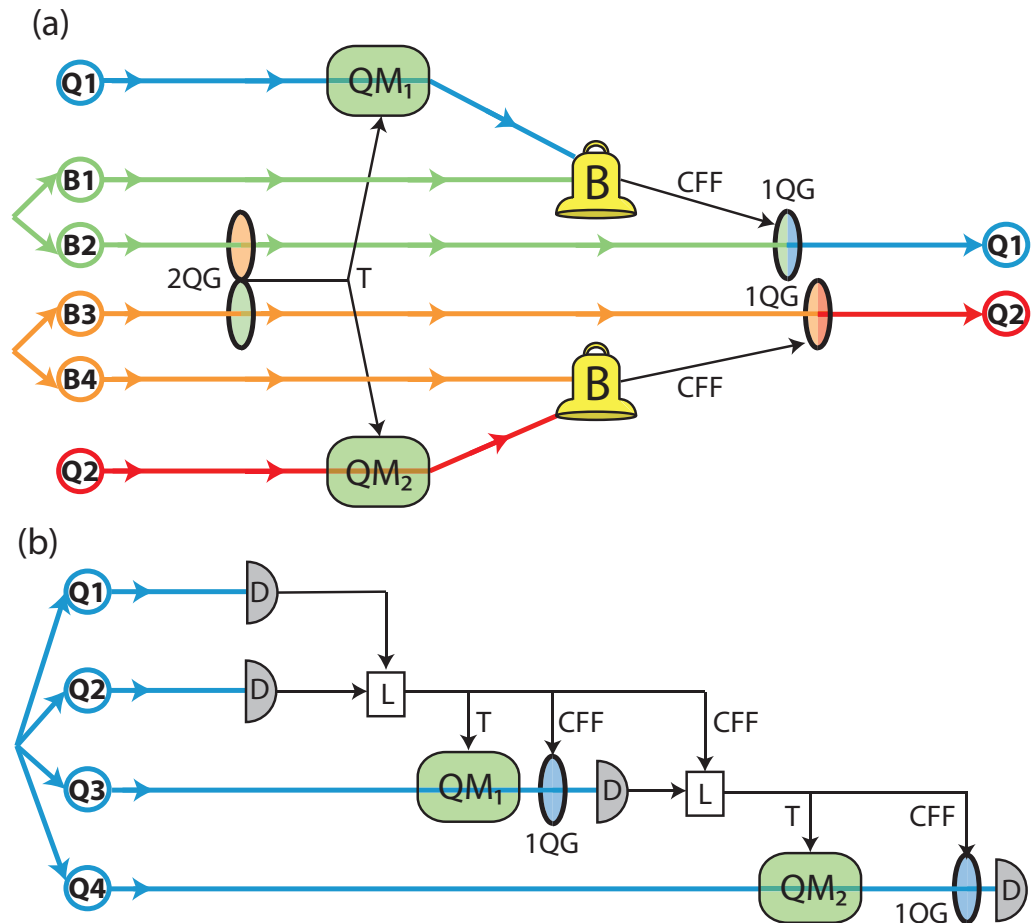


Figure 1.5: Quantum Computing. (a) Simplified Knill et al. protocol for linear-optic quantum computation [25]. One photon out of each of two Bell-state entangled pairs (B_2 and B_3) are used to attempt a two-qubit gate (2QG). At the same time the two qubits Q_1 and Q_2 are stored in quantum memories (QMs). Once the gate has been achieved, a trigger (T) is sent to the quantum memories and the states of Q_1 and Q_2 are teleported onto B_2 and B_3 by performing Bell-state measurements (B) and sending classical feedforward signals (CFF - black dashed lines) to one-qubit gates (1QG). (b) A four-qubit entangled state is used for one-way quantum computing via feedforward. Q_3 and Q_4 are stored in QMs to allow for synchronisation. A detection protocol (D) is performed on qubits Q_1 and Q_2 and logic (L) performed on the outcomes of these measurements is used to determine operations to apply to Q_3 before it, too, is detected. The measurement is also used to trigger QM_1 . The combined logic from detections Q_1 - Q_3 is used to alter Q_4 at another 1QG and trigger QM_2 . This is a simplified and modified version of the protocol presented in Ref. [56].

relevance to our daily lives, as most of the encrypted information we send and receive over public channels, such as the internet, is protected by the difficulty classical computers have in factoring large numbers [60].

A quantum computer needs to be able to perform single-qubit gates such as polarisation rotation, phase shifters and beam splitters, as well as two-qubit gates, namely a controlled-not gate [61]. There are many proposed protocols for a quantum computer, such as trapped ions [62, 63, 64] and solid state ensembles [65]. Here we will focus on linear-optic-based quantum computing (LOQC). The advantage of LOQC is that photons, as mentioned above, are great carriers of information, and performing single-qubit gates (such as polarisation rotations, beam-splitters and phase shifters) is straight-forward. Using photons would also allow for linking between LOQC and quantum communication channels. However, the fact that photons interact only very weakly with their environment means that performing two-qubit gates is quite difficult. These operations use the fact that, after approaching a beam-splitter, two indistinguishable photons will choose to both travel down the same path, experimentally demonstrated by C. K. Hong et al. in 1987 [66]. Unfortunately, with only these elements, deterministic gates cannot be formed and the system will be probabilistic. This will mean that it will need exponential resources, due to the nature of compounded probabilistic gates, and therefore offer no advantage over a classical computer.

There are two main methods for overcoming this. The first of these was proposed in 2001 by E. Knill et al. using detectors and teleportation [25]. Their idea was that each two-qubit gate is still probabilistic, but now heralded so that a successful gate is identified. Even so, for the times when the gates are unsuccessful, the two qubits would still be lost. To overcome this issue, they combined their idea with work by D. Gottesman et al. who, in 1999, proposed taking the gates “offline” and then teleporting them in [67]. This means that, though the gates will still fail with the same probability, the qubits of interest will not be destroyed and the non-unity probability of success will only reduce the gate creation rate.

An illustration of this scheme is shown in Fig. 1.5(a). Here two photons from two different Bell states (B2 and B3) are used to attempt the two-qubit gate, while the qubits Q1 and Q2 wait inside quantum memories until the gate has succeeded. Once success has been heralded, the states of Q1 and Q2 are teleported onto B2 and B3 by performing Bell-state measurements with their partner photons B1 and B4 and various single-qubit operations. Though this scheme greatly improves the possibility of success for a gate (for instance, for the controlled-phase gate described in Ref. [25] the total probability of success is $1/16$), it will still be non-unity, as with only linear-optics components we cannot completely distinguish between the four Bell states [68].

The other method of performing LOQC is with one-way, or cluster-state, quantum computing [69, 70, 71]. These protocols rely on entanglement, measurement and feed-forward techniques to process the initial state. In this situation quantum memories would be used as synchronisation devices so that feed-forward operations could be timed properly. For instance, in 2007 R. Prevedel et al. demonstrated fast LOQC using feed-forward, where fibre delay-lines were used to allow for feed-forward measurements [56]. This is illustrated in Fig. 1.5(b), modified to include quantum memories in place of the delay-lines. These one-way feed-forward schemes would be more robust if controllable delays were available via quantum memories, especially if more qubits were to be processed in parallel.

For more information on quantum computing, see review articles in Ref.s [72] and [73].




Figure 1.6: Single-Photon Gun. (i) Entangled photon pair generation and (ii) loading of the single-photon gun - a click on the detector (D) indicates that a photon has been stored in the quantum memory (QM). (iii) Firing the single-photon gun - by triggering (T) the QM we can retrieve a deterministic single-photon from the memory.

1.3.3 Single-Photon Sources

While not a quantum information process, deterministic single-photons are an important resource for both QKD (where extra photons can be used by Eve to attempt to remain undetected) and in quantum computing (where extra photons can lead to measurement and gate errors), as well as other processes such as quantum random number generation [18]. Many attempts have been made to produce an on-demand single-photon “gun” using single atoms [74] or ions [75] in cavities, quantum dots [76, 77] or nitrogen-vacancy (NV) centres in diamond [78]. However, all of these methods suffer from issues such as complex and expensive set-ups (single atoms/ions), the need for spectral filters that reduce efficiency (quantum dots) [79], or issues involving dark states (NV centres) [80].

Another option is to use heralded entangled photon pairs, such as those produced via spontaneous parametric down-conversion [81, 21], and a quantum memory. This is illustrated in Fig. 1.6. In this example, the detection of one photon indicates that the other is stored in the quantum memory and therefore we can turn the probabilistic process (with efficiencies of $10^{-7} \rightarrow 10^{-11}$ [80]) into a deterministic one, allowing us to recall the photon from the memory at any desired time.

For more information on single-photons, their production and uses, see Ref. [80].

1.3.4 Quantum Metrology

Finally, we will briefly discuss quantum metrology. This is a field where quantum phenomena such as squeezing and entanglement are used to allow for more precise measurements than are allowed for by classical theory. There is no better example of this than atomic frequency standards being used to define the second [82, 83]. Other areas where quantum processes can improve measurements include quantum lithography to beat the classical diffraction limit [84, 85], quantum positioning and clock synchronisation based on shared prior entanglement [86, 87] and ghost imaging where an image of an object is built (though it is never directly viewed) by using entangled photon pairs [88, 89].

Quantum memories can be useful for quantum metrology applications as their mapping of the quantum state from light to atoms means that measurements can be performed on stationary qubits, rather than ones travelling at the speed of light. For instance,

spin-squeezed states of atoms improve signal-to-noise for measurements in sensing and metrology [90]. Also, precision magnetometry can be implemented with entangled atoms [91].

For more information on quantum metrology see Ref. [3]. For more information on quantum images see Ref.s [92] and [93]. For examples of quantum memories improving metrology see Sec. 1.5.2.

1.4 Quantum Memory Metrics

Now that we have established the motivation for the development of a quantum memory we will discuss the necessary properties for a quantum memory to be useful for the tasks discussed above.

1.4.1 Efficiency

Efficiency is perhaps one of the easiest properties of a quantum memory to measure and understand. It is simply the ratio of output to input energy. For single-photon storage, the efficiency translates to the probability of photon re-emission from the memory. It is important to have high efficiencies ($> 90\%$) for single-photon states to allow for high key creation rates with quantum repeaters [55] as well as for deterministic single-photon sources. For multi-photon and continuous-variable state storage efficiency is also important, as a non-unit efficiency will lead to coupling of the vacuum into the quantum state, leading to its degradation. Care must be taken when measuring the efficiency as other, detrimental, effects such as amplification of the state may occur, leading to misleadingly good results.

Efficiency is normally linked to how strongly the photons interacts with the storage medium. This can be characterised by the optical depth (OD) of a memory. The OD is not the whole story, as different techniques require more or less OD to achieve the same efficiencies. Efficiencies down to 50% can still be useful for applications such as quantum computing if error correction protocols are used [94].

1.4.2 Fidelity

The fidelity of a quantum memory is a measure of the overlap of the recalled state with the input state. This is illustrated in Fig. 1.7 for a Wigner function state representation [95]. Measuring the fidelity is one way of determining whether noise has been added to the output state. Below 50% fidelity recall, a classical measure-and-reproduce memory could work just as well [96]. If we can achieve fidelities over 68% for coherent states we pass the no-cloning limit [97]. This is where the output is guaranteed to be the best possible copy of the input state. For single-photons we can decouple efficiency and fidelity by looking at the conditional fidelity, i.e., only using measurement outcomes where a photon is detected.

One issue with fidelity is that it is state dependent. As can be seen from Fig. 1.7, for a constant non-unit efficiency, overlap of two uncertainty “balls” (determining the fidelity) will be different depending on the size of the original amplitude “stick”. Perhaps a more accurate, state-independent, measure of how well the output state matches the input state is using a T-V diagram [98, 99].

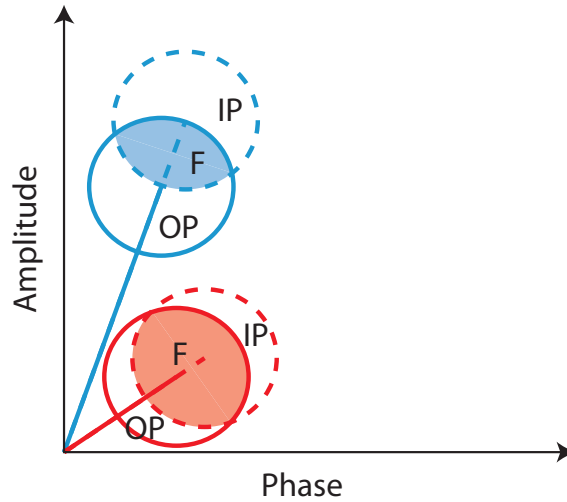


Figure 1.7: Fidelity. Wigner function ball-on-stick state representation for two coherent states (see next chapter) with different amplitudes: IP - input state (dashed); OP - output state (solid); F - fidelity (shaded regions). Both output states have same percentage reduction of stick length.

Both fidelity and T-V diagrams will be discussed in more detail in Sec. 3.2.1.

1.4.3 Storage Time

The coherence time of a memory is a measure of how quickly the stored information is lost, normally measured as the $1/e$ decay point. Even a quantum memory that has an efficiency approaching unity for short storage times will be useless if the information decays too quickly for its intended application. For quantum repeaters, for instance, we require a coherence time much longer than the time it takes information to travel between several nodes [100]. This is normally on the order of milliseconds. As an upper bound, we would ideally want information to last within a quantum repeater until final entanglement distillation is complete. This can be up to hundreds of seconds [55]. We also want to be able to control the storage time and be able to recall information from the memory as quickly as possible after triggering the release.

1.4.4 Bandwidth

The bandwidth of a memory will determine how narrow a pulse can be stored and therefore limit the repetition and bit rates for a system. The required bandwidth will depend on the application as well as the source. For instance, single photons generated by down-conversion normally have bandwidths on the order of nanometers [101], though using cavities these can be reduced to under ten megahertz [102, 103].

1.4.5 Delay-Bandwidth Product

Apart from the absolute value of coherence times and bandwidths of a memory, another important metric is the delay-bandwidth product (DBP). Simply put, this is the ratio of storage time to pulse length. To be useful we want a large DBP to ensure the coherence time is much longer than the input pulse duration. The DBP also determines the number of pulses, or bits of information, that could be stored in the memory at one time. However, a memory may have a large DBP while not being able to store multiple modes (see next section).

1.4.6 Multi-Mode Capacity

The multi-mode capacity of a memory determines how many different modes can be stored in the memory at one time. These could be different spatial [104], temporal [105] or frequency modes. Apart from improving bit rates, the multi-mode capacity of a memory can also be used to help improve quantum repeater protocols [105, 106]. For instance, with a multi-mode memory, the required storage time for full entanglement distillation reduces from hundreds of seconds down to seconds.

1.4.7 Wavelength

Another property to consider for a quantum memory is the wavelength at which it operates. For long-distance communication we would ideally want a memory that functioned at around 1500 nm, the lowest absorption band for optical fibres.

1.4.8 Storage Media

Though not technically a metric of performance, different storage media lend themselves towards certain applications over others. Normally ensembles of atoms are used as the storage medium, though single atoms can also be used (see Sec. 1.5.1). These are usually contained either in solid state ensembles cooled down to cryogenic temperatures, or alkali atoms, either in warm vapour cells or cooled and trapped using lasers. Memories using nitrogen-vacancy centres in diamond are also being touted for potential future memories at room temperature.

When choosing a storage medium we want one that has a high OD to allow for high efficiencies. The storage medium will have a large impact on the coherence times that can be achieved. For instance, in some media atoms are free to move, limiting coherence times due to diffusion and inter-atomic collisions, while others hold atoms rigidly in place. The storage medium will also determine which wavelengths can be used. In the end, the choice of storage medium is nearly as important as the memory scheme used for satisfying the criteria listed above, as will be seen in the next section.

1.5 A Trip to the Quantum Memory Zoo

Having defined the properties we desire in a quantum memory, we will now view some of the main quantum memory candidates. As will be noticed by the reader, these candidates are composed of a vast array of different beasts. Like a trip to any zoo, there is not enough time to view all the animals, but we will endeavour to highlight the main genera and a few of the more exotic varieties. We will only present species that have been experimentally sighted, avoiding the mythical beasts that have been theoretically proposed but have of yet not been demonstrated.

1.5.1 Delay Lines and Cavities

Perhaps the simplest organism in the quantum memory world is a loop of fibre. The advantages of using such a delay line are that it is readily available, easy to use and has a high bandwidth. However, it does not allow for controllable storage times. Even long delay times are not possible as fibres have a minimum absorption of 0.2 dB km^{-1} at 1500 nm [45] and therefore 50% of the light will be lost after 15 km, a storage time of roughly $70 \mu\text{s}$. For storage times on the order of a millisecond, efficiencies will be effectively zero. Despite these issues, fibres have been used to delay entanglement for $125 \mu\text{s}$ [107], and theoretical proposals made to improve their performance using feedback and error correction [108].

A slightly more evolved version of the fibre loop was demonstrated by T. B. Pittman et al. in 2002 where they used a free-space loop containing dynamic polarisation control to be able to recall 700 nm single-photons at multiples of the initial cycle time (13 ns), with approximately 20% loss per cycle [109]. This led to a quasi-coherence time of 50 ns.

Taking this idea a step further leads to cavity-based quantum memories. The idea is very similar to a fibre loop, with light being coupled into the cavity and then left to bounce around inside for a set amount of time. The number of cycles the light will undergo inside a cavity will depend on its quality-factor (Q). Cavities, as with fibre-loops, suffer from fixed storage times and there are limitations on the modes that satisfy resonance conditions. However, if the Q can be controlled, then storage times can be altered, as was the case in the experiment of T. Tanabe et al., where they could delay an approximately 1 ns pulse by 1.45 ns and then tune the Q of their photonic crystal cavity ($\approx 10^6$) with pump pulses [110]. Because of the small round-trip distance in such a cavity (on the order of micrometers), very large finesse will be needed for useful storage times. And, while Qs up to $4 \cdot 10^{10}$ have been achieved with microtoroids [111], there will be a trade-off between finesse and coupling efficiency [112].

However, if we were to put an atom inside a cavity then we have a different story. In fact, some of the first proposals for quantum memories relied on the strong coupling that can be achieved between light and atoms when they are surrounded by a cavity [113]. J. I. Cirac et al. proposed using atoms in cavities as nodes in a quantum network with entangled single-photons in 1997 [53] and it was soon after, in 1999, that the first experiment demonstrating the localisation of an atom in a cavity was performed by J. Ye et al. [114]. Since that time, atoms have been localised inside a cavity for over ten seconds [115, 116]. Using this method, H. P. Specht et al. were able to demonstrate storage with 9.3% efficiency (limited by atom-cavity coupling), with a coherence time of $180 \mu\text{s}$ (measured from the fidelity) for single-photon storage at 780 nm using a ^{87}Rb atom [117]. However, this technique is very difficult to implement, and will only work at the few-photon level.

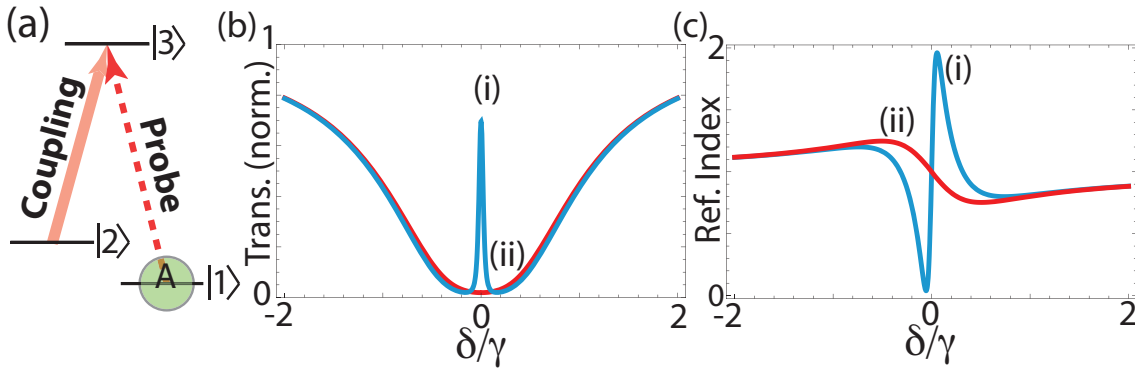


Figure 1.8: Electromagnetically-Induced Transparency (a) Level scheme for EIT, showing probe (dashed) and coupling field (solid) transitions, as well as atomic population (A). (b) Transmission profile and (c) the refractive index for an EIT medium as a function of probe-only detuning (normalised to the excited-state decay rate) both (i) with (blue) and (ii) without (red) a coupling field.

1.5.2 Electromagnetically-Induced Transparency

Electromagnetically-induced transparency (EIT) is one of the most studied, and therefore most developed, quantum memory candidates. The basic EIT scheme is illustrated in Fig. 1.8(a). As can be seen, it is based around an ensemble of quantised three-level atoms, with ground states $|1\rangle$ and $|2\rangle$ and an excited state $|3\rangle$. If we were to send in our quantum signal encoded on a weak (probe) pulse of light resonant with the $|1\rangle \rightarrow |3\rangle$ transition then, if we had a large number of atoms present, the whole quantum signal would be absorbed. This would not be useful as we would have no way of then retrieving the excitation. However, if we were to also shine a strong (coupling) light field resonant with the $|2\rangle \rightarrow |3\rangle$ transition onto the atoms, we would create a transparency window around the probe transition. This is shown in Fig. 1.8(b), and the corresponding change in refractive index is shown in Fig. 1.8(c). The theory behind this effect will be discussed in more detail in Ch. 3.

As can be seen from Fig. 1.8(c), with the coupling field present there is a sharp change in refractive index around the transmission window for the probe and this will lead to a corresponding reduction in the group velocity of the information. Using EIT, velocities down to the order of tens of meters per second have been achieved [118, 119]. By slowly reducing the amount of coupling light present we can reduce the width of the transparency window and therefore decrease the group velocity further. If we adiabatically reduce the coupling field power to zero, it follows that the probe pulse will become trapped inside the atomic cloud. This was the idea for EIT as a quantum memory proposed by M. Fleischhauer et al. in 2000 [120]. We can characterise this as a stopped-light effect, with the information being stored in a coherence between the two ground states.

The first experimental demonstrations of classical light storage with EIT were by C. Liu et al. and D. F. Phillips et al. in 2001, using cold atomic ensembles [121] and warm vapour cells [122] respectively. The first demonstration using a solid state system was also performed in the same year by A. V. Turukhin [123]. Since then efficiencies in EIT have




Figure 1.9: Raman Memory. (a) The level scheme for Raman memory, showing probe (dashed) and coupling field (solid) transitions and atomic population (A). (b)(i) Write operation, where probe and coupling field are both pulsed to create a coherence between the two ground states. (ii) Read operation, where a secondary coupling field pulse recalls the coherence. Coupling field leakage through the quantum memory (QM) not shown.

reached up to 43% with warm vapours [124] and 78% with cold atomic ensembles [13]; both of these made use of optimisation techniques [125] and using light between 780-795 nm. The latter experiment had a coherence time of 98 μs and a DBP of 74 for greater than 50% efficiency. A coherence time of 2.3 s for EIT was demonstrated by J. J. Longdell et al. in 2005 using rephasing pulses in a solid state ensemble with light at approximately 600 nm [5].

Apart from storage of classical light, EIT has also been used to store single-photons [126, 127], as well as the delay [128] and storage [129, 130, 131] of squeezed vacuum pulses. Entanglement has been stored using a pair of EIT memories by K. S. Choi et al. in 2008 [132]. EIT has also been used in quantum metrology as a frequency standard [133, 134] and in high precision magnetometry [135].

As well as achieving these remarkable results, there are a few issues associated with light storage using EIT. Firstly, it has been found both theoretically [136, 137] and experimentally [138] that EIT adds noise to the signal, most probably due to coupling-field-induced scattering [139]. There are also issues with achieving large bandwidths in EIT, as a pulse must both spectrally fit within the EIT transmission window and spatially fit within the memory length. The first condition sets a minimum pulse length (maximum bandwidth) and the second condition sets a maximum pulse length (minimum bandwidth). This trade-off also limits the multi-mode ability for EIT, with it scaling with the square root of the OD [140]. To achieve high bandwidths and allow for multi-mode storage therefore requires large ODs ($\gg 1$). Also, with high ODs and a resonant Λ -system, four-wave mixing becomes an issue. This will be discussed more in Sec. 5.3.5.

1.5.3 Raman Memory

Raman memory, first proposed by J. Nunn et al. in 2007 [141], uses a Λ -system and additional coupling field like EIT but, unlike EIT, this time both fields are detuned from resonance. This is illustrated in Fig. 1.9. By being off-resonance the behaviour of the system without the coupling field is completely different, in that the probe will pass

straight through the memory. The presence of the coupling field allows a coherence to be built between the two ground states (i.e., a spin-wave). Storing information in the ground states of a real atom is normally desirable as excited states have relatively short coherence times due to spontaneous emission. The information can also only be recalled with the application of a second coupling pulse, making it a controllable process.

The input pulse will have a non-zero bandwidth, inversely proportional to its temporal width. Therefore, to allow the full pulse to be absorbed when there is a small inhomogeneous linewidth, the “write” coupling field must also be pulsed. If the two are mode-matched then a large storage bandwidth can be achieved. The information is then retrieved with a second “read” coupling pulse.

This technique was first demonstrated using a warm vapour cell in 2010, when K. F. Reim et al. were able to store a 300 ps pulse, with a bandwidth of approximately 1.5 GHz [7]. This work was extended further, with single-photon-level storage with efficiencies of 30% and a coherence time of 1.2 μ s [142]. This gave them a DBP of 2,500. Both experiments used light at 850 nm. It was also found that, with the off-resonant nature of the memory, a beam-splitter operation could be performed on the stored information [143].

The efficiencies achieved were limited by lack of coupling field power to increase the effective OD, as well as reabsorption issues for retrieval in the forwards direction (this will be discussed more in Sec. 1.5.5). Though a large DBP can be achieved with the Raman memory, it is effectively single-mode in its simplest form [144]. This is because, to write another pulse into the memory requires turning on the coupling field, and that will lead to recall of the already-stored information. Noise measurements performed on the Raman scheme showed that coupling-field-induced scattering added thermal photons to the detection [142].

1.5.4 The Duan-Lukin-Cirac-Zoller Protocol

The Duan-Lukin-Cirac-Zoller (DLCZ) protocol, named after its architects, is designed to create long-lived, long-distance entanglement between memories, with quantum repeater applications in mind [1]. As with the previous two schemes presented, the DLCZ protocol uses an ensemble of Λ -type atoms and, like the Raman memory, it is detuned from resonance. The basic scheme is illustrated in Fig. 1.10(a). The system is prepared by sending a weak off-resonance “write” pulse, near the $|1\rangle \rightarrow |3\rangle$ transition, into the memory. This will probabilistically create a spin excitation, whereby one atom is moved from the $|1\rangle$ state to state $|2\rangle$, heralded by the emission of a Stokes photon. To retrieve the excitation, a “read” pulse, off-resonance from the $|2\rangle \rightarrow |3\rangle$ transition, is sent into the memory to create an anti-Stokes photon.

To create entanglement between different memories, beam-splitters are used to combine the anti-Stokes output ports from two memories, and read pulses applied to both. In this situation, if there is a photon detected we will know that one memory has fired but not know which, leaving them entangled. This is illustrated in Fig. 1.10(b). One advantage of this protocol over other quantum-memory-based repeater protocols (i.e., see Fig. 1.4) is that entangled photon sources are not required; instead detectors and beam-splitters should be placed between the memories. Also entanglement purification is built into the scheme [1].

First demonstrated in 2003 for both many-photons in a warm vapour [145] and for




Figure 1.10: DLCZ Scheme. (a) The DLCZ level scheme with (i) writing operation where a write pulse (red, solid) detuned from resonance is sent into an ensemble of atoms (A) causing the excitation of a Stokes photon (orange, dashed), the detection of which indicates the memory is “loaded.” (ii) The reading operation where a read pulse (orange, solid) detuned from resonance is sent into the ensemble to cause the emission of an anti-Stokes photon (red, dashed). (b) Entangling two quantum memories (QMs) by firstly loading them and then applying read pulses. The outputs from the QMs are combined on a beam-splitter (BS) so that a detection event (D) indicates that one memory has fired its anti-Stokes photon, but without knowing which one.

single-photons in a cold atomic ensemble [146], retrieval efficiencies of up to 50% for 30-ns-long pulses stored for 300 ns have been demonstrated without a cavity, using cold atomic ensembles at 850 nm [147]. With a cavity to define the mode of the anti-Stokes photons, this has been increased to 84% for another cold atomic ensemble with a coherence time on the order of hundreds of nanoseconds also at 850 nm [148]. Another cavity arrangement with cold atoms has led to a slightly lower recall efficiency of 73%, but a much longer coherence time of 3.2 ms at 795 nm [149].

As can be seen, the DLCZ protocol lends itself to producing deterministic single-photon sources. It can also be used to create pairs of entangled photons [150]. Moving towards quantum repeater applications, entanglement between two remote atomic ensembles has been demonstrated [151], and even four at once [152]. One issue with the DLCZ protocol is that it is single-mode and care must be taken to ensure only one atom is excited during the write process.

1.5.5 Photon Echo Techniques

The idea for a photon echo memory stemmed from the concept of nuclear-spin echoes proposed by E. L. Hahn in 1950 [153] and was first implemented by N. A. Kurnit et al. in 1964 using a ruby crystal [154]. Photon echo memories are based around the idea of reversible absorption from an ensemble of atoms. If we imagine, as illustrated in Fig. 1.11, an ensemble of two-level atoms initially in the ground state, then all their dipole vectors will be aligned downwards. If we were to excite the ensemble with enough light to put it in an equal superposition of its ground and excited states, called a $\pi/2$ -pulse, then the dipoles should still be aligned if the absorption was coherent. However, due to the inhomogeneous broadening of the ensemble, each atom will then start to rotate, or precess, around the Bloch sphere with a frequency determined by its detuning from resonance. If this were all




Figure 1.11: Two-Level Photon Echo Scheme. (a) The level structure (include atomic population (A)) and (b) Bloch sphere representation of the photon echo process. (i) A $\pi/2$ -pulse places the ensemble in an equal superposition of ground and excited states. (ii) Each dipole then rotates around the Bloch sphere with a frequency determined by its detuning from the pulse centre frequency (rotating frame of reference) (iii) A π -pulse inverts the medium causing (iv) rephasing of the dipoles and therefore emission of a photon echo.

we were to do then the excitation will remain trapped inside the ensemble until it decayed via spontaneous emission. However, if we were to apply another pulse that rotated the dipole vectors by 180° (i.e., inverted the population), called a π -pulse, then the dipoles will realign and, by a reversal of the absorption process, build a macroscopic coherence to emit a coherent pulse of light - a photon echo.

There are, however, major concerns about using this two-pulse photon echo for quantum applications. Namely, the π -pulse that is used to invert the population can lead to spontaneous emission from the ensemble and therefore amplification and added noise [155]. This is especially an issue for storing single-photons, which we have seen are so useful for quantum information processing. There is also a reabsorption issue as, if the ensemble has a high optical depth, then light attempting to leave the memory will have a high likelihood of being absorbed again. This leads to a maximum retrieval efficiency of 54% in the forwards direction [156], and is an issue for most transmissive memories. However, 100% recall can be achieved in the backwards direction.

As with many organisms facing extinction, photon echoes have had to evolve. We will discuss some of these variants below. However, before this we should note that the extra noise generated by the two-pulse photon echo can be useful for quantum repeater applications. In 2010, P. M. Ledingham et al. proposed a scheme that would use the extra noise to create a DLCZ-type quantum repeater [157]. This technique was named rephased amplified spontaneous emission (RASE), where the extra noise is rephased and a stream of photons are stored and recalled, rather than just one. This was first implemented in 2012 by S. E. Beavan et al. in a solid state ensemble [158]. Around the same time, P. M. Ledingham et al. were able to demonstrate correlations between the amplified spontaneous emission and RASE fields [159]. RASE has an advantage over the DLCZ protocol in terms of its multi-mode ability and large bandwidth.




Figure 1.12: Three-Pulse Photon Echo Scheme. A $\pi/2$ write pulse prepares the ensemble for the input pulse. A $\pi/2$ read pulse is sent in a time τ after the input to retrieve an echo at time 2τ .

Three-Pulse Photon Echo

In 1982 T. M. Mossberg et al. described a way to store arbitrary pulses using a three-pulse photon echo technique [160]. Their idea being that a write pulse, with area $\pi/2$ and a bandwidth as large as the information to be stored, is sent into the ensemble first to prepare the memory. The information is then sent in and, after a time τ , another $\pi/2$ -pulse is sent in. As with the π -pulse in the two-level photon echo scheme, this will cause rephasing of the dipoles and the emission of a photon echo. This is illustrated in Fig. 1.12.

Using this technique, storage has been demonstrated for up to 3.8 hours in a solid state system at 580 nm [161], up to 4000 pulses have been stored at one time in another solid state ensemble at 790 nm [162] and signal processing at gigahertz rates has been achieved using yet another solid state system, also at approximately 790 nm [163]. However, this scheme suffers from the same amplification and noise issues as the two-level scheme, meaning that high efficiencies cannot be achieved without population inversion and the corresponding spontaneous emission. Modified photon echo schemes have produced published efficiencies of 20% in 2012 [164].

In 2011, two groups proposed and demonstrated a photon-echo technique that works without these drawbacks by suppressing the first echo and then rephasing it. D. L. McAuslan et al. used electric fields to suppress the first echo intensity by over 99% in a solid state ensemble at 606 nm [165]. The efficiency for the final echo, however, was only a few percent in this case. V. Damon et al. used spatial phase mismatch to suppress the first echo [166]. They named this technique ROSE - rephasing of silenced echo - and demonstrated a recall efficiency of 10% after storage for 25 μs in a solid state ensemble at 790 nm.

Atomic Frequency Combs

Atomic Frequency Combs (AFCs), proposed by M. Afzelius et al. [167], are another way of rephasing the atomic dipoles without the noise and amplification issues discussed above. In its simplest form, AFC uses an ensemble of two-level atoms that is prepared with an inhomogeneous broadening resembling a comb pattern, with a set spacing δ_c between each comb. This is shown in Fig. 1.13(a). When we then send a pulse to excite the storage medium, each dipole will precess with a speed that is a multiple of the other dipoles around it. This means that, after a time $t_s = 2\pi/\delta_c$, all the dipoles will realign and therefore emit a photon echo in the forwards direction.




Figure 1.13: AFC Scheme. (a) Comb structure of AFC absorption window in frequency space, with δ_c being the frequency spacing between combs. (b) Delay-line AFC storage. After the absorption of the probe (P) the dipoles precess at a rate proportional to their detuning (i.e., multiples of δ_c). This means that at $t = 2\pi/\delta_c$ they will realign to emit an echo (E) from the ensemble. (c) Controllable storage level structure for AFC, where coupling pulses are used to transfer the coherence to and from an auxiliary state $|Aux\rangle$.

The preparation of a storage medium with a comb structure has meant that, to date, only solid state ensembles have been used for AFC memories. The first demonstration of AFC was with weak coherent pulses by H. de Riedmatten et al. in 2008 [168] and since then efficiencies up to 35% have been achieved at 605 nm [169]. One of the main advantages of the AFC scheme is that it is highly multi-mode, scaling with the number of teeth in the comb (which, in turn, is limited by the inhomogeneous broadening and the excited state decay rate) [167]. With this multi-mode capacity, and bandwidths up to 5 GHz at 795 nm [170], 1060 pulses have been stored at once [8]. Entanglement storage with true single-photons has been demonstrated with AFCs [171], as has entanglement between two memories [172] and storage of time-bin qubits [6].

However, there are some issues with this two-level AFC scheme. Firstly, recall in the forward direction is still limited to 54%, due to reabsorption, and backwards retrieval is not possible. Secondly, the storage time is fixed so that it is more a delay-line than a quantum memory. Finally, storing light in an excited state will limit storage times due to the higher excited-state decay rate. All these issues can be overcome by adding an auxiliary ground-state $|aux\rangle$ and using coupling pulses to move the excitation from the upper state to the auxiliary state before the rephasing time t_s and, by using counter-propagating pulses, the echo could be recalled backwards. Three-level storage was first experimentally demonstrated by M. Afzelius et al. in 2010, where they showed controllable storage with a maximum efficiency of 1% with two modes and a coherence time of approximately 20 μs for recall in the forwards direction at approximately 600 nm [173]. Very recent results have improved on this to have 6% efficiency and demonstrate storage of five modes, again at 600 nm [174].

Controlled Reversible Inhomogeneous Broadening

Another way of producing an echo is via controlled reversible inhomogeneous broadening (CRIB). The idea here is that, if the inhomogeneous broadening of each atom could be somehow reversed, then we could time-reverse the initial absorption process exactly to produce a photon echo with high efficiency. The original proposal, by S. A. Moiseev et al. in 2001 [175], envisaged a warm vapour medium and used the fact that the Doppler




Figure 1.14: CRIB Scheme. (a) The level structure of the atomic ensemble (A) and pulses applied (and times) for the CRIB scheme. The coherence between the states $|1\rangle$ and $|3\rangle$ created by a probe (P) pulse is transferred to state $|2\rangle$ with a π -pulse and then transferred back with a counter-propagating π -pulse (assuming a Doppler-broadened ensemble) to cause re-emission of an echo (E) in the backwards direction. (b) Bloch sphere representation of absorption and emission. (i) Input, with dipoles precessing at a frequency determined by their Doppler-shifted detuning. (ii) When coherence is transferred back to the $|1\rangle \rightarrow |3\rangle$ transition the detunings will be reversed due to nature of the Doppler shift, causing realignment of the dipoles.

shift of light seen by a moving atom (in the rest frame of the lab) is opposite for counter-propagating lasers. If we were to use a Λ -style atomic ensemble again, and send in the information to be stored on the $|1\rangle \rightarrow |3\rangle$ transition, a pair of counter-propagating π -pulses would allow for storage and retrieval, respectively, from state $|2\rangle$. As long as the time between the input pulse and initial π -pulse is much smaller than the spontaneous emission rate from state $|3\rangle$, then the opposite Doppler shifts due to the two π -pulses will reverse the direction of the dipole precession on the $|1\rangle \rightarrow |3\rangle$ transition and cause an echo to be emitted. This is illustrated in Fig. 1.14.

This idea was then extended to solid state systems (where atoms have no velocity) by controlling the inhomogeneous broadening of the atoms externally, for instance with electric or magnetic fields [176, 177]. The counter-propagating π -pulses were still required to allow for up to 100% recall efficiency in the backwards direction. It was found that a two-level implementation without the π -pulses was also possible, but that retrieval in the forwards direction was limited to 54% [156]. In terms of multi-mode ability, it was found that the number of modes scales linearly with the OD, better than EIT but lower than AFC [140].

The first demonstration of this two-level, frequency-gradient-reversal technique was by A. L. Alexander et al. in 2006 in a solid state ensemble using reversible electric fields for rephasing. They improved on their results a year later with up to four pulses stored, as well as demonstrating conservation of phase inside the memory, but with efficiencies on the order of 10^{-5} at 580 nm [178]. In 2010, B. Lauritzen et al. were able to use the same method with photons at telecom wavelengths (around 1550 nm), again in a solid state ensemble [179]. They were able to demonstrate storage with 0.2% efficiency after approximately 300 ns.

Gradient Echo Memory

From a taxonomy standpoint the gradient echo memory (GEM) scheme should be classed as longitudinal CRIB. In any case, as it is the main focus of this thesis we will discuss it separately. GEM follows directly from the realisation that only a two-level atom and a reversible frequency gradient are required for the emission of a photon echo. However, it was the discovery that a *monotonic* gradient placed along the memory could lead to 100% recall efficiency in the forwards direction that sets it apart. This was first put forward by G. Hétet et al. in 2008 when they performed another gradient-reversing echo experiment in a solid state ensemble with an efficiency of 15% after approximately 3 μ s at 600 nm [180]. This reversal mechanism is illustrated in Fig. 1.15.

The frequency gradient is really the key to GEM, not just causing the rephasing of the atomic dipoles, but the bandwidth of the memory is determined by the product of the gradient η and the length of the memory. The monotonicity of the gradient ensure that the recalled light will not be absorbed as it leaves the ensemble and therefore allows for the 100% forward recall efficiency. It also makes GEM a frequency-encoding memory, as different frequencies of the pulse will be stored at different spatial locations along the memory (i.e., z -direction).

Much experimental work has been undertaken on GEM since these first experiments. In its initial incarnation, i.e., two-level solid state scheme, efficiencies up to 69% have been achieved, with no added noise to place its performance above the no-cloning limit at 600 nm [181]. This was achieved by M. P. Hedges et al. in 2010. The bandwidth of the

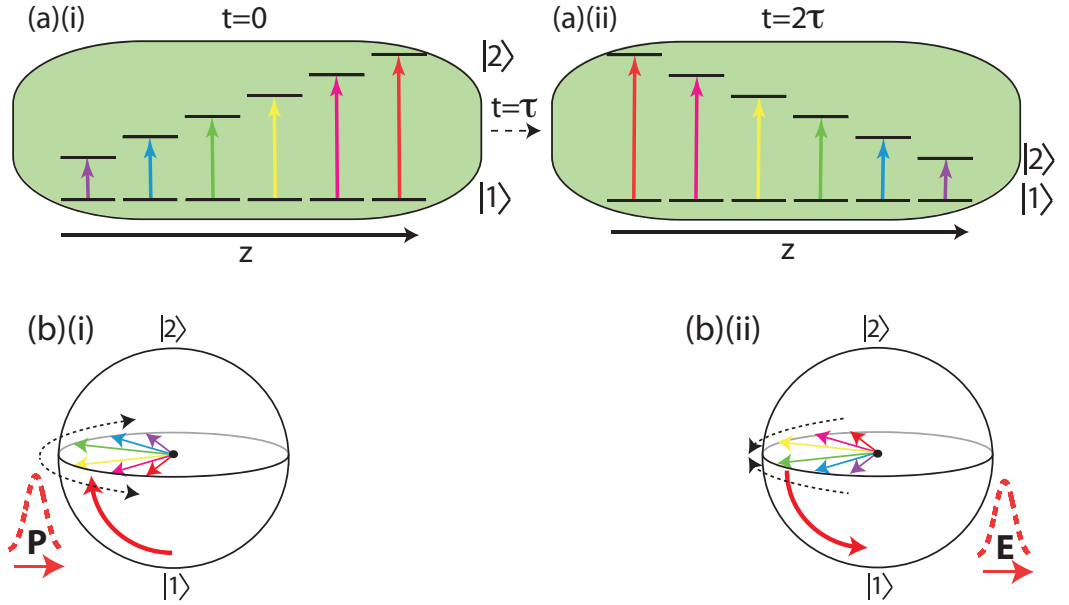


Figure 1.15: The GEM Scheme. (a) Level structure and (b) Bloch diagram of GEM scheme. (i) Pulse enters ensemble with a frequency gradient η in the z -direction. Dipoles precess with a frequency determined by the gradient-induced detuning at their location. (ii) At time $t = \tau$ the gradient is reversed ($\eta \rightarrow -\eta$), reversing the direction but not speed of the dipoles. This causes re-emission from the memory at time $t = 2\tau$ in the forwards direction.

memory was approximately 1 MHz.

A three-level version of GEM, known as Λ -GEM, has also been developed using a warm rubidium-87 vapour cell [182, 183]. The initial reason for the change was the convenience of using vapour cells, in terms of ease of use, laser wavelengths (795 nm), and the high ODs that can be achieved. The change in level structure was to avoid using the excited state of ^{87}Rb that would otherwise limit the memory lifetime to tens of nanoseconds [184]. The detuned three-level structure with the addition of a coupling field, similar to that shown in Fig. 1.9(a), also allows for decoupling of the OD and coherence time on the same parameter. This quasi-two-level arrangement should, in principle, allow for long storage times with high efficiencies. Efficiencies up to 87% have been demonstrated so far, making Λ -GEM the most efficient quantum memory to date [4]. Again the bandwidth was on the order of 1 MHz. Up to twenty pulses have also been stored using Λ -GEM and a coherence time of approximately 20 μs measured in the same experiment.

Performing operations on stored information using both GEM and Λ -GEM have also been theoretically proposed, such as manipulation using the gradient [185]. Simple operations such as stretching or compressing a pulse by recalling with a steeper or shallower gradient have already been demonstrated [9]. A frequency offset has also been added [4]. On top of this, the extra “knob” for Λ -GEM in terms of the coupling field allows us the ability to efficiently resequence pulses inside the memory in Λ -GEM [9], as well as perform controllable beam-splitter operations [186, 10]. Other operations such as using GEM as an optical router [187] or all-optical multiplexer [188] have also been put forward, demonstrating the flexibility of the scheme.

Fully quantum states have not yet been stored using GEM, though tomography

was performed on two-level GEM using weak coherent states to characterise the noise properties. One issue with storing states such as single-photons in GEM is the relatively small bandwidth of roughly one megahertz. The current limitation on the bandwidth is OD, as the effective broadened OD scales as $1/\text{bandwidth}$. Also, in Λ -GEM the effective two-level OD is much lower than the OD of the $|1\rangle \rightarrow |3\rangle$ transition, meaning that high ODs are needed for efficient storage. The multi-mode capability of GEM scales linearly with the OD, the same as CRIB.

1.5.6 Other Species

Apart from the “big four” above (EIT, Raman, AFC and CRIB/GEM), there are many other quantum memories being developed. Here we will take a look at a select few to give a flavour of what is out there.

Four-Wave Mixing

Four-wave mixing (FWM) can be both a very useful process, as well as a detrimental one in other quantum memory protocols (see Sec. 5.3.5). FWM has been used to generate entanglement (see, for example, Ref.s [189] and [190]) and squeezing [191, 192] from warm atomic vapours at wavelengths and bandwidths compatible with quantum memories. Not many experiments, however, have used FWM alone to store information. FWM takes place in a Λ -system similar to EIT, where the coupling field can also interact with the $|1\rangle \rightarrow |3\rangle$ transition and produce an extra light field. As with EIT, delay of information has been demonstrated [193], most notably with entangled images [194]. Storage has also been demonstrated in a warm vapour cell, with efficiencies of a few percent and a coherence time of $120 \mu\text{s}$ at 795 nm [195].

Off-Resonance Faraday Interactions

The first demonstration of mapping quantum properties of light into an ensemble of atoms was performed by B. Julsgaard et al. in 2004 [196]. The authors were able to map the polarisation state of photons onto atoms using two vapour cells, a polarisation-sensitive detection scheme, and feedback onto the atoms by applying a magnetic field. It was shown that the fidelity of the writing operation was 70%, therefore above the classical limit, for up to 4 ms at 850 nm. Using a similar method of measurement and feedback, K. Jensen et al. demonstrated storage of continuous-variable squeezed-states in warm vapour cells with a lifetime of 1 ms [197]. However, retrieving the information from the memory in these cases could prove difficult, potentially requiring some form of atom-light teleportation (i.e., see Ref. [198]).

Stimulated Brillouin Scattering

If we were to pass a high intensity laser through an optical fibre it would cause the fibre to change shape due to the non-uniform intensity distribution. This will create phonons than can then interact with other photons. By sending in strong write and read pulses, counter-propagating with the information pulses, information can be controllably stored and retrieved using these phonons. With this method, efficiencies up to 29% have been

	Efficiency	Fidelity	Coherence Time
Value	87% [4]	97.5% (av., cond.)[202]	2.3 s [5]
Scheme	Λ -GEM	AFC	EIT
Medium	Warm Vapour	Solid State	Solid State
Wavelength	795 nm	883 nm	606 nm
Λ -GEM Best [4]	87%	Not yet measured	$\approx 20 \mu\text{s}$
	Bandwidth	DBP	Multi-Mode Storage
Value	5 GHz [170]	2,500 [142] (74 > 50% [13])	1,060 [8]
Scheme	AFC	Raman (EIT)	AFC
Medium	Solid State	Warm Vapour (Cold Atoms)	Solid State
Wavelength	795 nm	850 nm (780 nm)	793 nm
Λ -GEM Best [4]	≈ 1 MHz	≈ 40 (11 > 50%)	20

Table 1.1: Quantum Memory Performance Summary. N.B. the value of 40 for DBP for Λ -GEM stems from the fact that in the 20-pulse storage experiment the first pulse sent in was the last pulse recalled. Bracketed terms in DBP refer to the DBP with recall above the classical 50% efficiency limit. av. - averaged, cond. - conditional.

achieved for approximately 2 ns pulses with a coherence time on the order of nanoseconds [199]. One benefit of this method is that it allows for storage at telecom wavelengths, as fibres are used as the storage medium. They were also able to store up to three pulses in the memory at one time.

Nuclear Spin Memory

Nuclear spins are ideal for storage of information as they are much more isolated than electron spins and therefore decay at a much slower rate. Using a single nuclear spin of a ^{13}C atom in a diamond nitrogen-vacancy centre at room temperature, P. C. Maurer et al. were able to demonstrate a coherence time of over one second [200]. This was achieved using a green laser field and a sequence of radio-frequency pulses. Published at the same time, M. Steger used a ^{28}Si atom's nuclear spin to show a coherence time of three minutes at 2 K [201]. One issue with using the nuclear spin is the low electron-to-nuclear spin coupling.

1.6 Conclusions and Thesis Outline

I hope that by the end of this chapter the reader is convinced of the importance of developing a quantum memory as it is the key piece of quantum hardware at the heart of long-distance provably-secure communication, while at the same time being able to help break current encryption methods and having the potential to allow us to measure quantities better than ever before. I also hope that the reader can appreciate that there are many different schemes out there, each with their own strengths and weaknesses. Tab. 1.1 gives a summary of the best results achieved for the metrics discussed in Sec. 1.4. As can be seen, if we were able to combine the best of these together we would already have developed a useful quantum memory.

Sadly, not all these properties are available in one memory and so we must do our best to find a candidate that satisfies as many metrics as possible and attempt to improve it. This is the goal of my thesis, using Λ -GEM. The best performance metrics for Λ -GEM are also shown Tab. 1.1. As can be seen, Λ -GEM is a promising candidate due to the high efficiencies achieved.

As such, the experiments presented in my thesis are split into two main parts: in Ch.s 5-8 we investigate ways of improving the functionality of Λ -GEM in warm vapour cells, from digitising the set-up and performing noise measurements, to spectral manipulation experiments and demonstrating Λ -GEM's ability to store multiple spatial modes. Then in Ch.s 9-12 we propose and implement a move from warm to cold atoms to allow for alternate gradient creation techniques, and attempt to improve Λ -GEM's performance in terms of longer storage times while maintaining high efficiencies, we also take a slight detour to investigate cold atom storage with EIT. Before all of this, we present the development of digital control software for quantum optics experiments in Ch. 4, and conclude the introductory part of the thesis by presenting all relevant theory in Ch.s 2 and 3. This outline is illustrated in Fig. 1.16.

The majority of this thesis has been published, submitted to, or accepted for publication in international journals. Some selected articles resulting from the work done during my Ph.D. and included in this thesis are:

- B. M. Sparkes, M. Hosseini, G. Hétet, P. K. Lam, and B. C. Buchler, *ac Stark gradient echo memory in cold atoms*, Physical Review A **82**, 043847 (2010)
- B. C. Buchler, M. Hosseini, G. Hétet, B. M. Sparkes, J. J. Longdell, M. J. Sellars and P. K. Lam, *High Efficiency Gradient Echo Memory with 3-Level Atoms*, AIP Conference Proceedings **1363**, 383 (2010)
- M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler, *A Room Temperature Quantum Memory*, Nature Physics **7**, 794 (2011)
- H. M. Chrzanowski, J. Bernu, B. M. Sparkes, B. Hage, A. P. Lund, T. C. Ralph, P. K. Lam, and T. Symul, *Photon-number discrimination without a photon counter and its applications to reconstructing non-Gaussian states*, Physical Review A **84**, 050302(R) (2011)
- B. M. Sparkes, H. M. Chrzanowski, D. P. Parrain, B. C. Buchler, P. K. Lam, and T. Symul, *A Scalable, Self-Analyzing Digital Locking System for use on Quantum Optics Experiments*, Review of Scientific Instruments **82**, 075113 (2011)
- M. Hosseini, B. M. Sparkes, G. Campbell, B. C. Buchler, and P. K. Lam, *A Room Temperature Quantum Optical Memory*, International Conference on Quantum Information, OSA Technical Digest (CD), paper QTuG1 (2011)
- B. M. Sparkes, M. Hosseini, G. Hétet, P. K. Lam, and B. C. Buchler, *Spectral Manipulation of Optical Pulses Using the Gradient Echo Memory Scheme*, Proceedings of the International Quantum Electronics Conference and Conference on Lasers and Electro-Optics Pacific Rim 2011, paper I427 (2011)
- D. Higginbottom, B. M. Sparkes, M. Rancic, O. Pinel, M. Hosseini, P. K. Lam, and B. C. Buchler, *Spatial-mode storage in a gradient-echo memory*, Physical Review A **86**, 023801 (2012)

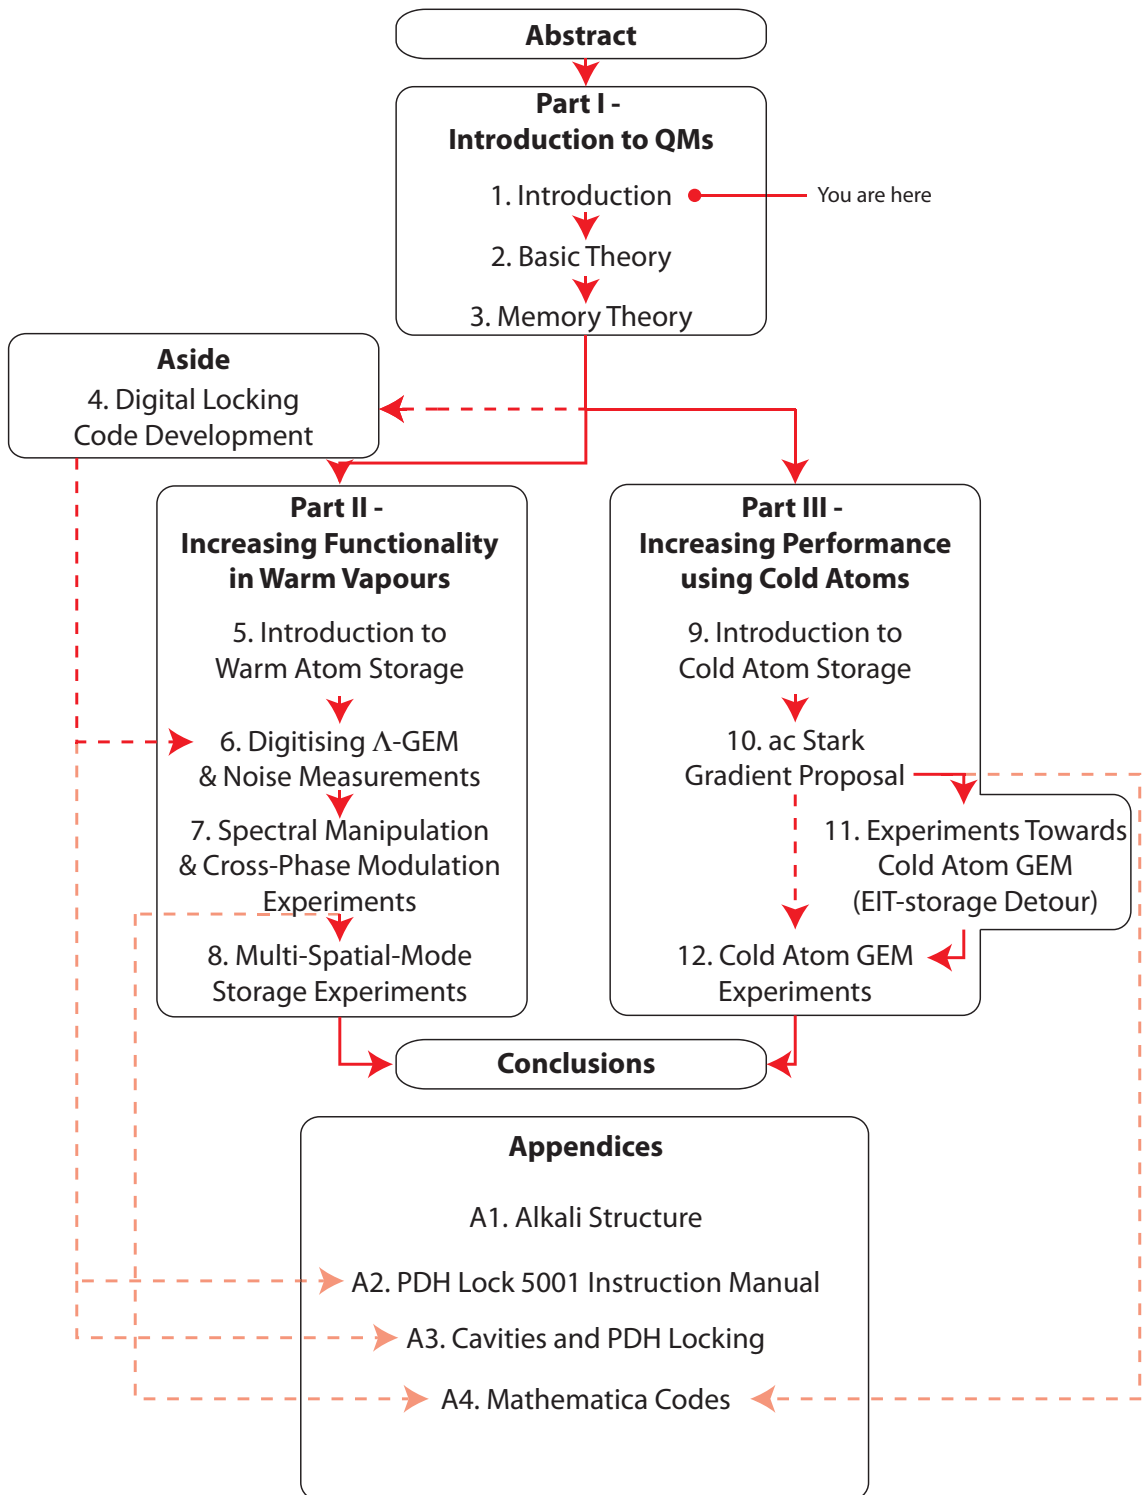


Figure 1.16: Thesis Outline

- B. M. Sparkes, C. Cairns, M. Hosseini, D. Higginbottom, G. Campbell, P. K. Lam, and B. C. Buchler, *Precision spectral manipulation: a demonstration using a coherent optical memory*, *Physical Review X* **2**, 021011 (2012)
- M. Hosseini, S. Rebić, B. M. Sparkes, J. Twamley, B. C. Buchler, and P. K. Lam, *Memory-enhanced noiseless cross-phase modulation*, *Light:Science & Applications* **1**, e40 (2012)
- L. Giner, L. Veissier, B. M. Sparkes, A. S. Sheremet, A. Nicolas, O. S. Mishina, M. Scherman, S. Burks, I. Shomroni, D. V. Kupriyanov, P. K. Lam, E. Giacobino, and J. Laurat, *Experimental Investigation of the Transition between Autler-Townes Splitting and Electromagnetically-Induced Transparency*, *Physical Review A* **87**, 013823 (2013)
- B. M. Sparkes, J. Bernu, M. Hosseini, J. Geng, Q. Glorieux, P. A. Altin, P. K. Lam, N. P. Robins, and B. C. Buchler, *Gradient echo memory in an ultra-high optical depth cold atomic ensemble*, arXiv:1211.7171v3 (2012).

Other articles published during my Ph.D. include:

- M. Hosseini, B. M. Sparkes, G. Hetet, J. J. Longdell, P. K. Lam, and B. C. Buchler, *Coherent optical pulse sequencer for quantum applications*, *Nature* **461**, 241 (2009).
- 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).
- 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 (2010).
- G. Campbell, M. Hosseini, B. M. Sparkes, P. K. Lam, and B. C. Buchler, *Time- and frequency-domain polariton interference*, *New Journal of Physics* **14**, 033022 (2012).
- M. Hosseini, B. M. Sparkes, G. T. Campbell, P. K. Lam, and B. C. Buchler, *Storage and manipulation of light using a Raman gradient echo process*, *Journal of Physics B* **45**, 124004 (2012).

Quantum Theory of Atom-Light Interactions

Sheldon: “I’m a physicist, I have a working knowledge of the entire universe and everything it contains.” - “The Work Song Nanocluster” The Big Bang Theory, CBS, 16 March 2009, TV Series.

2.1 Introduction

In this chapter we will develop the quantum mechanical formalism required to understand the workings of electromagnetically-induced transparency and the gradient echo memory, which will be presented in the next chapter. We will also add mathematical scaffolding to better understand some of the issues related to quantum information processing presented in the previous chapter. We will start by introducing quantum mechanical formalism in Sec. 2.2, before discussing different quantum states of light (Sec. 2.3), atoms (Sec. 2.4) and their interactions (Sec. 2.5).

This is not designed to be a comprehensive theoretical introduction to quantum atom-optics. Interested readers looking for more detail may wish to read Ref.s [203, 204, 205, 206, 207, 208, 209] as well as the thesis of G. Hétet [210].

2.2 Quantum Formalism

This section details the relevant formalism needed to understand the issues associated with measurement in quantum mechanics, as well as the time-evolution of states.

2.2.1 Quantum States and Operators

In quantum mechanics we can represent the system using state vectors $|\psi\rangle$ (in Dirac’s bra-ket notation). These are complex vectors that live within a complex vector space called the Hilbert space, the dimensions of the Hilbert space being the number of degrees of freedom of the property described by the state (or equivalently the number of orthogonal basis vectors). An inner product $\langle\psi'|\psi\rangle$ is defined on this Hilbert space. We can represent an arbitrary state in the Hilbert space as a sum or orthonormal basis states $\{|\psi_j\rangle\}$, i.e., $|\psi\rangle = \sum_j c_j |\psi_j\rangle$ where c_j are complex numbers, $|c_j|^2$ is the probability of finding the system in state $|\psi_j\rangle$ and $\sum_j |c_j|^2 = 1$ for normalisation. We call this a coherent superposition of the

basis states. We can also define a state as being a statistical mixture of the basis states with probability $|c_j|^2$, i.e., $|\psi\rangle = \{|c_j|^2, |\psi_j\rangle\}$.

The difference between a coherent superposition and a statistical mixture may not be readily apparent. However, if we define a density matrix $\hat{\rho}$ to be

$$\hat{\rho} = |\psi\rangle \langle\psi| = \sum_{j,k} c_{jk} |\psi_j\rangle \langle\psi_k|, \quad (2.1)$$

where $c_{jk} = c_j \cdot c_k^*$ - the star denoting the complex conjugate. We can now differentiate between the two because, though both have populations $c_{jj} = |c_j|^2$, a statistical mixture will have $c_{jk} = 0$ for all $j \neq k$, while this is not the case for a superposition state. We call the c_{jk} terms the coherences, and they determine the phase between the states. This means that a coherent superposition is simultaneously in multiple states at once, and we can see interference between these (think, for example, of a single photon going through two slits and interfering with itself). A statistical mixture, meanwhile, is either in one state or another and, without any coherence terms, no interference can occur. We can therefore see how quantum-mechanical phenomena are linked to these coherences, as mentioned in the introduction.

Two important properties of the density matrix are

$$\text{Tr}\{\hat{\rho}\} = \sum_j \langle\phi_j|\hat{\rho}|\phi_j\rangle = 1 \quad (2.2)$$

$$\text{Tr}\{\hat{\rho}^2\} = \sum_j \langle\phi_j|\hat{\rho}^2|\phi_j\rangle \leq 1, \quad (2.3)$$

where the trace (Tr) is over an arbitrary basis $\{|\phi_j\rangle\}$. If the trace in Eq. 2.3 equals one, then we have a pure state where the system is in one of the eigenstates $|\phi_j\rangle$ with unit probability. We can also define operators \hat{O} that act on states in the Hilbert space as follows

$$\hat{O} = \sum_{j,k} o_{jk} |\phi_j\rangle \langle\phi_k|. \quad (2.4)$$

If $o_{jk} = o_{kj}^*$ then $\hat{O} = \hat{O}^\dagger$. We call these operators Hermitian.

2.2.2 Quantum Measurements and Uncertainty

Before measuring a quantum system we first have to know what properties we can observe. We can associate these observables with a set of Hermitian measurement operators $\{\hat{M}_j\}$, where j refers to the measurement outcome, i.e.,

$$\hat{M} = \sum_j m_j \hat{M}_j = \sum_j m_j |\psi_j\rangle \langle\psi_j|, \quad (2.5)$$

where $\{m_j\}$ are a set of real eigenvalues corresponding to the measurement eigenstates $\{|\psi_j\rangle\}$. The average, or expectation value, for the measurement will be given by

$$\langle\hat{M}\rangle = \text{Tr}\{\hat{M}\hat{\rho}\} = \sum_j m_j c_{jj}, \quad (2.6)$$

for the state described by the density matrix in Eq. 2.1. After a measurement with result j the system will no longer be in a superposition of the measurement states, but solely in the state $|\psi_j\rangle$. This is known as the collapse of the wavefunction, and illustrates how performing a measurement on a quantum state will alter it. This will be the case except for a pure state (in the measurement basis).

Apart from wavefunction collapse, there is also the issue of inherent uncertainty in a quantum state. For instance, if we wanted to obtain information about two observables, represented by measurement operators \hat{M} and \hat{M}' , we cannot measure both to an arbitrary degree of certainty if the following is true

$$[\hat{M}, \hat{M}'] = \hat{M}\hat{M}' - \hat{M}'\hat{M} \neq 0. \quad (2.7)$$

If this is the case then we say the two observables are non-commuting. This is encapsulated in the Heisenberg uncertainty principle (HUP), which states that for two non-commuting observables the following relation holds

$$\Delta\hat{M}\Delta\hat{M}' \geq \frac{\hbar}{2}, \quad (2.8)$$

where $\Delta\hat{M}$ is the uncertainty in measurement of \hat{M} , mathematically given by

$$\Delta\hat{M} = \sqrt{\langle\hat{M}^2\rangle - \langle\hat{M}\rangle^2}. \quad (2.9)$$

It is now obvious why a classical measure-and-reproduce memory is not adequate for the uses presented in the previous chapter: firstly performing a measurement leads to the collapse of the state; and secondly there is a finite limit on how much we can know about a state through measurement.

2.2.3 Pictures and Time-Evolution

Perhaps the most important operator of all is the Hamiltonian $\hat{\mathcal{H}}$. This is an Hermitian operator that is used to determine the total energy E of a system via

$$\hat{\mathcal{H}}|\psi\rangle = E|\psi\rangle. \quad (2.10)$$

It also determines the time evolution of the system. In this thesis we are mostly concerned with the evolution of operators and will therefore use the Heisenberg picture where

$$\frac{d}{dt}\hat{O} = \frac{1}{i\hbar}[\hat{O}, \hat{\mathcal{H}}] + \frac{\partial}{\partial t}\hat{O}. \quad (2.11)$$

The partial derivative in the above equation will equal zero if the operator does not contain explicit time-dependence. We can also calculate the evolution of the density matrix via the master equation

$$\frac{\partial\hat{\rho}}{\partial t} = -\frac{i}{\hbar}[\hat{\mathcal{H}}, \hat{\rho}] + \mathcal{L}[\hat{\rho}], \quad (2.12)$$

where $\mathcal{L}[\hat{\rho}]$ is the Liouvillian of the system, describing the effects of coupling the system to an external reservoir.

To make calculations as simple as possible, we can split the Hamiltonian into two parts: $\hat{\mathcal{H}} = \hat{\mathcal{H}}_o + \hat{\mathcal{V}}$, where $\hat{\mathcal{H}}_o = \hat{\mathcal{H}}_l + \hat{\mathcal{H}}_a$ is the Hamiltonian of the individual systems

(in our case the light and the atoms respectively) and $\hat{\mathcal{V}}$ is the interaction Hamiltonian, describing their interactions. In the remainder of this chapter we will discuss the forms of $\hat{\mathcal{H}}_l$, $\hat{\mathcal{H}}_a$ and $\hat{\mathcal{V}}$.

2.3 Quantum Theory of Light

In this section we will present the quantum mechanical description of light, focusing mainly on those states that we will use later in this thesis.

2.3.1 Number States, Creation and Annihilation Operators

The building blocks of quantum states of light are photons, and the easiest way of representing a state with a definite number of photons is, perhaps unsurprisingly, the number state basis. In this basis, a state with n photons in mode k (with angular frequency $\omega_k = c/k$, c being the speed of light) is described by $|n_k\rangle$. Two important operators that act on number states in mode k are the creation (\hat{a}_k^\dagger) and annihilation (\hat{a}_k) operators, following the rules

$$\hat{a}_k^\dagger |n_k\rangle = \sqrt{n_k + 1} |n_k + 1\rangle \quad (2.13)$$

$$\hat{a}_k |n_k\rangle = \sqrt{n_k} |n_k - 1\rangle \quad (2.14)$$

$$[\hat{a}_k, \hat{a}_j^\dagger] = \delta_{kj}, \quad (2.15)$$

where δ_{kj} is the Kronecker delta ($\delta_{kk} = 1$, $\delta_{kj} = 0$ for all $k \neq j$). These operators are non-Hermitian and therefore do not represent observables of the system. We can, however, create a very useful observable from them. This is the number operator $\hat{N}_k = \hat{a}_k^\dagger \hat{a}_k$ that acts on a number state, without altering it, to determine the number of photons present. I.e.,

$$\hat{N}_k |n_k\rangle = n_k |n_k\rangle. \quad (2.16)$$

The Hamiltonian for a single-mode number state, taking the energy of the vacuum state ($|0\rangle$) to have zero energy (as we are only interested in the change of energy between states), gives

$$\hat{\mathcal{H}}_l = \hbar\omega_k \hat{N}_k. \quad (2.17)$$

2.3.2 Indistinguishability, Entangled States and Teleportation

Single-photons in the same mode are indistinguishable and this can lead to some interesting results. For instance, if we image two single photons, indistinguishable apart from their position denoted by the modes a and b , approaching a 50:50 beam-splitter at the same time (as shown in Fig. 2.1) then we can determine their paths after the beam splitter to be

$$|1_a\rangle \rightarrow \frac{1}{\sqrt{2}} (|1_c\rangle + |1_d\rangle) \quad (2.18)$$

$$|1_b\rangle \rightarrow \frac{1}{\sqrt{2}} (|1_c\rangle - |1_d\rangle), \quad (2.19)$$

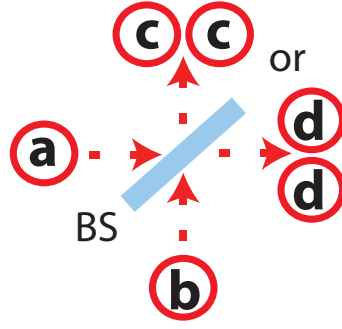


Figure 2.1: Indistinguishability. Two identical photons in modes a and b interfere on a 50:50 beam-splitter (BS), with the outcome being that both leave in either mode c or mode d but not one in each.

where the negative sign for the second mode comes from the need to conserve energy. Now if we look at the total state (ignoring normalisation) we have

$$\begin{aligned} |1_a\rangle |1_b\rangle &\rightarrow (|1_c\rangle + |1_d\rangle)(|1_c\rangle - |1_d\rangle) \\ &= (|2_c\rangle - |2_d\rangle + |1_c\rangle |1_d\rangle - |1_d\rangle |1_c\rangle), \end{aligned} \quad (2.20)$$

$$(2.21)$$

and as modes c and d are orthogonal and the photons indistinguishable we can commute then to remove the last two terms and produce the state

$$\frac{1}{\sqrt{2}} (|2_c\rangle - |2_d\rangle). \quad (2.22)$$

This property of photons was discussed in the previous chapter as a way of creating two-photon gates for a linear-optic-based quantum computer.

Another useful two-photon state is an entangled state, this is where we have correlated photon pairs. Mathematically we can determine if a two-photon state is entangled or not by seeing if we can factorise the combined state. For instance, if we were to use the tried-and-true example of two photons, in modes a and b , in either horizontal ($|H\rangle$) or vertical ($|V\rangle$) polarisation-states, then the following state would be entangled

$$|\Phi_{ab}^+\rangle = \frac{1}{\sqrt{2}} (|H_a\rangle |H_b\rangle + |V_a\rangle |V_b\rangle). \quad (2.23)$$

Another way of viewing this is that, if we were to perform a polarisation measurement on the photon in mode a and find it to be horizontal, then we know that the photon in mode b must also be horizontal. This is how Alice and Bob can generate their key from an entangled photon pair in the Ekert quantum key distribution protocol discussed in Sec. 1.3.1.

The state shown above is one of four maximally-entangled Bell states, the others are given below

$$|\Phi_{ab}^-\rangle = \frac{1}{\sqrt{2}} (|H_a\rangle |H_b\rangle - |V_a\rangle |V_b\rangle) \quad (2.24)$$

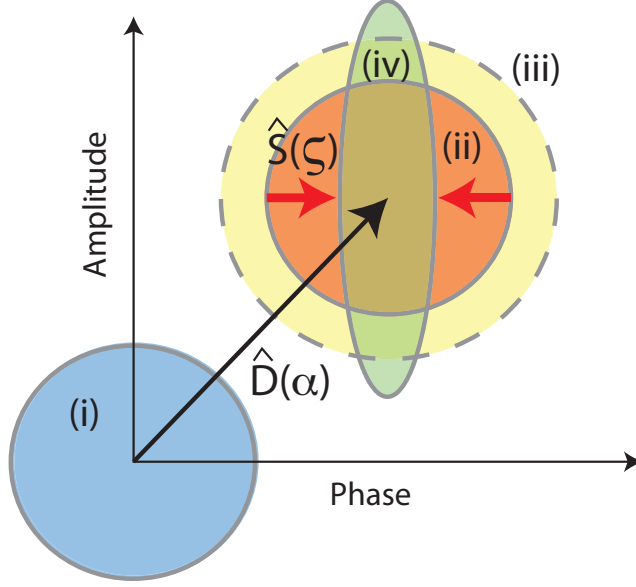


Figure 2.2: Coherent, Squeezed and Thermal States. Wigner function ball-on-stick representations of (i) vacuum state (blue), (ii) coherent state (red), (iii) thermal state (yellow, dashed line), and (iv) squeezed state (green).

$$|\Psi_{ab}^{\pm}\rangle = \frac{1}{\sqrt{2}} (|H_a\rangle |V_b\rangle \pm |V_a\rangle |H_b\rangle). \quad (2.25)$$

These states can also be used to teleport information. For instance, imagine we had a qubit in state $|\psi_q\rangle = c_H |H_q\rangle + c_V |V_q\rangle$, assuming normalisation. If we were to combine this qubit with one photon from a Bell-state pair, let's take $|\Psi^{-}\rangle$, then the state of the whole system will be given by

$$\begin{aligned} |\psi_{abq}\rangle &= \frac{1}{\sqrt{2}} (c_H |H_q\rangle + c_V |V_q\rangle) (|H_a\rangle |V_b\rangle \pm |V_a\rangle |H_b\rangle) \\ &= \frac{1}{2} (|\Phi_{qa}^{+}\rangle (c_H |V_b\rangle - c_V |H_b\rangle) + |\Phi_{qa}^{-}\rangle (c_H |V_b\rangle + c_V |H_b\rangle) \\ &\quad + |\Psi_{qa}^{+}\rangle (-c_V |H_b\rangle + c_H |V_b\rangle) + |\Psi_{qa}^{-}\rangle (c_V |H_b\rangle + c_H |V_b\rangle)). \end{aligned} \quad (2.26)$$

From the above equation we can see that if we are able to distinguish between the four Bell-states then we will know what operation (i.e., phase flip on one polarisation state, or rotation of the total polarisation by $\pi/2$) is required on the photon in mode b to make it indistinguishable from the initial qubit. A measurement that can determine the exact Bell-state of the entangled pair of photons is, unsurprisingly, called a Bell-state measurement. It can be seen from Eq. 2.23-2.25 above that using only polarising beam-splitters we can determine between two of the four in any given measurement (i.e., both photons through one port, or one through each).

2.3.3 Coherent and Squeezed States

Fixed-number photons states are difficult to produce experimentally and we will therefore look slightly further afield to describe other, more easily producible, quantum states. The most common form of a single-mode light states is the coherent state. These states are eigenstates of the annihilation operator and can be described in terms of number states to be

$$|\alpha_k\rangle = \text{Exp}[-|\alpha|^2/2] \sum_{n=0}^{\infty} \frac{\alpha^n}{(n!)^{1/2}} |n_k\rangle. \quad (2.27)$$

From this definition we can quickly determine that the mean number of photons in a coherent state is given by $\langle \hat{N} \rangle = |\alpha|^2$. For $\alpha = 0$ the coherent state and number state representations are the same. Indeed, we can build coherent states from this vacuum state ($|0\rangle$) using the displacement operator $\hat{D}_k(\alpha) = \text{Exp}[\alpha \hat{a}_k^\dagger - \alpha^* \hat{a}_k]$ as follows

$$\hat{D}_k(\alpha) |0\rangle = |\alpha_k\rangle. \quad (2.28)$$

A state similar to a coherent state is a thermal state. However, a thermal state is a statistical mixture of number states of the form [206]

$$\hat{\rho}_{th} = \frac{1}{\mathcal{G}} \left(|0\rangle \langle 0| + \frac{\mathcal{G}-1}{\mathcal{G}} |1\rangle \langle 1| + \left(\frac{\mathcal{G}-1}{\mathcal{G}} \right)^2 |2\rangle \langle 2| + \dots \right), \quad (2.29)$$

where \mathcal{G} is a parameter dependent on the temperature of the thermal field, becoming larger as the temperature increases. The other difference between a coherent and a thermal state is that a coherent state follows Poissonian statistics for photon number, while a thermal state has a super-Poissonian distribution.

For both coherent and thermal states we can define amplitude and phase operators in terms of the creation and annihilation operators. They are respectively defined as

$$\hat{X}_{1,k} = \frac{1}{2} (\hat{a}_k + \hat{a}_k^\dagger) \quad (2.30)$$

$$\hat{X}_{2,k} = \frac{-i}{2} (\hat{a}_k - \hat{a}_k^\dagger). \quad (2.31)$$

Using the commutation relation for \hat{a}_k and \hat{a}_k^\dagger from Eq. 2.15 we can determine the commutator for these two observables to be

$$[\hat{X}_{1,k}, \hat{X}_{2,j}] = \frac{i}{2} \delta_{jk}. \quad (2.32)$$

We therefore know, by the HUP, that there will be an uncertainty relation between the amplitude and phase “quadratures” for a coherent state. For a coherent state we can determine that the uncertainty in both quadratures is the same, that is $\Delta \hat{X}_{1,k} = \Delta \hat{X}_{2,k}$, and both are at the Heisenberg limit. This is true regardless of the value of α , even for the vacuum state. For a thermal state the uncertainties are also equal, but larger than the Heisenberg limit. These are shown in Fig. 2.2. Coherent states, as pure states with minimum uncertainty, are useful for applications such as quantum tomography (discussed in Sec. 6.3) as well as for the generation of Schrödinger cat states (discussed more in Sec. 4.5) where a superposition of two of these “quasi-classical” states $|\alpha\rangle$ and $|\alpha\rangle$ can be thought of as the dead-alive superposition of Schrödinger’s ill-fated pet.

As stated in the introduction, we can reduce the uncertainty in one of these quadrates if we are prepared for a corresponding increase in the other. This is achieved via a squeezing operator $\hat{S}_k(\zeta) = \text{Exp} \left[\frac{1}{2}\zeta^* \hat{a}_k^2 - \frac{1}{2}\zeta \hat{a}_k^{\dagger 2} \right]$, where ζ is the squeezing parameter. As we will not be working with squeezing in this thesis we will not discuss it in detail except to say that the squeezed state is then given by

$$|\alpha_k, \zeta\rangle = \hat{D}_k(\alpha) \hat{S}_k(\zeta) |0_k\rangle, \quad (2.33)$$

and this is again illustrated in Fig. 2.2.

2.3.4 The Quantised Multi-Mode Electric Field

In the above sections we have only considered single-mode fields. We will now expand this to a continuum of fields. Firstly, solving the Maxwell equations in the Coulomb gauge gives an equation for the quantised, multi-mode electric field in terms of the creation and annihilation operators¹

$$\hat{E}(\mathbf{r}, t) = i \int_{-\infty}^{\infty} dk \sqrt{\frac{\hbar\omega_k}{2V\epsilon_o}} \left[\hat{a}_k u_k(\mathbf{r}) \text{Exp}[-i\omega_k t] - \hat{a}_k^\dagger u_k^*(\mathbf{r}) \text{Exp}[i\omega_k t] \right], \quad (2.34)$$

where $\mathbf{k} = (\omega_k/c)\mathbf{r}$ is the wave-vector in the direction of the unit vector \mathbf{r} , c being the speed of light in a vacuum, ϵ_o is the permittivity in a vacuum, $V = L^3$ is the quantisation volume for a box of length L , and $u_k(\mathbf{r}) = \text{Exp}[i\mathbf{k} \cdot \mathbf{r}]$ are the unnormalised orthogonal transverse spatial modes. We will discuss spatial modes in more detail below.

We can simplify the above equation if we assume that there are only a small range of frequencies $\delta\omega$ present in our light field, centred around $\omega_o = c|\mathbf{k}_o|$, and that $L \ll 1/|\mathbf{k}|$. This allows us to define a multi-mode slowly-varying envelope operator

$$\hat{\mathcal{E}}(\mathbf{r}, t) = i \int_{-\infty}^{\infty} dk' \hat{a}_{k'} u_{k'}(\mathbf{r}). \quad (2.35)$$

The electric field operator then becomes

$$\hat{E}(\mathbf{r}, t) = \sqrt{\frac{\hbar\omega_o}{2\epsilon_o V}} \left(\hat{\mathcal{E}}(\mathbf{r}, t) \text{Exp}[-i\omega_o t] - \hat{\mathcal{E}}^\dagger(\mathbf{r}, t) \text{Exp}[+i\omega_o t] \right). \quad (2.36)$$

The commutation relation for $\hat{\mathcal{E}}(\mathbf{r}, t)$ is given by

$$\left[\hat{\mathcal{E}}(\mathbf{r}, t), \hat{\mathcal{E}}^\dagger(\mathbf{r}', t') \right] = L\delta(ct - \mathbf{r} - (ct' - \mathbf{r}')), \quad (2.37)$$

where $\delta(t - t_o)$ is the Dirac delta function (which equals zero unless $t = t_o$ in which case it equals one) and the L comes from the normalisation of the $u_k(\mathbf{r})$ modes. We can also define new multi-mode amplitude and phase operators to be, respectively,

$$\hat{X}^+(\mathbf{r}, t) = \frac{1}{2} \left(\hat{\mathcal{E}}(\mathbf{r}, t) + \hat{\mathcal{E}}^\dagger(\mathbf{r}, t) \right) \quad (2.38)$$

$$\hat{X}^-(\mathbf{r}, t) = \frac{-i}{2} \left(\hat{\mathcal{E}}(\mathbf{r}, t) - \hat{\mathcal{E}}^\dagger(\mathbf{r}, t) \right). \quad (2.39)$$

The Hamiltonian for this multi-mode field will be given by an extension of the single-mode

¹For a complete derivation see, for Example, Ref. [208].




Figure 2.3: Transverse Electromagnetic Modes. Contour plots of the absolute value of the amplitude for different TEM modes: (i) TEM_{00} ; (ii) TEM_{01} ; (iii) TEM_{11} ; (iv) TEM_{02} ; (v) TEM_{12} ; (vi) TEM_{03} . Signs indicate the relative phases of the peaks, colour bar indicates relative peak heights.

case from Eq. 2.17 to be

$$\hat{\mathcal{H}}_l = \hbar \sum_k \omega_k \hat{a}_k^\dagger \hat{a}_k. \quad (2.40)$$

We will now return to discuss the different spatial modes $u_k(\mathbf{r}) = \text{Exp}[i\mathbf{k} \cdot \mathbf{r}]$. These must satisfy the following equation for a wave travelling in the z -direction and bounded in a volume V , following from classical electromagnetism

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} - 2ik \frac{\partial}{\partial z} \right) u_k(\mathbf{r}) = 0. \quad (2.41)$$

One solution to this equation is the Hermite-Gauss transverse electromagnetic modes (TEM). These are of the form [211]

$$\begin{aligned} \text{TEM}_{mn}(x, y, z) &= \sqrt{\frac{2}{\pi w(z)^2}} \frac{1}{\sqrt{m!n!2^{m+n}}} H_m \left(\frac{\sqrt{2}x}{w(z)} \right) H_n \left(\frac{\sqrt{2}y}{w(z)} \right) \text{Exp} \left[-\frac{x^2 + y^2}{w(z)^2} \right] \\ &\times \text{Exp} \left[-ik \frac{x^2 + y^2}{2R(z)} + i(m+n+1)\phi_G(z) - ikz \right], \end{aligned} \quad (2.42)$$

where the transverse beam radius, radius of curvature of the beam wave-front, Rayleigh length and Gouy phase shift are given by, respectively,

$$w(z) = w_o \sqrt{1 + (z/z_r)^2} \quad (2.43)$$

$$R(z) = z \left(1 + \frac{z_r^2}{z^2} \right) \quad (2.44)$$




Figure 2.4: Heterodyne Detection Set-Up. The weak probe field \mathcal{E} is combined with the strong local oscillator (LO) on a 50:50 beam-splitter (BS). The intensity signals measured at detectors D1 and D2 are subtracted to produce the heterodyne signal, supplying both amplitude and phase information about the probe.

$$z_R = \frac{\pi w_o^2}{\lambda} \quad (2.45)$$

$$\phi_G(z) = \text{Tan}^{-1} \left[\frac{z}{z_R} \right], \quad (2.46)$$

where w_o is the waist of the laser beam, with wavelength $\lambda = 2\pi/k$. The $H_{m,n}$ terms are the Hermite polynomials of the form

$$H_m(x) = (-1)^n \text{Exp} [x^2] \frac{d^n}{dx^n} \text{Exp} [-x^2]. \quad (2.47)$$

Some of these modes are plotted in Fig. 2.3.

2.3.5 Detection

As most current detectors cannot detect signals oscillating at ω_o , we will use the multi-mode slowly-varying envelope operator to define detection. As we defined the number operator to determine how many photons there were in a mode, we can similarly define a photo-current operator

$$\hat{i}(t) = \hat{\mathcal{E}}(t)^\dagger \hat{\mathcal{E}}(t). \quad (2.48)$$

Now, for a pulse of information we can simplify $\hat{\mathcal{E}}_p(t)$ by splitting it into its amplitude and phase quadrature components, $\hat{\mathcal{E}}_p^+(t)$ and $\hat{\mathcal{E}}_p^-(t)$ respectively. The photocurrent will therefore be

$$\begin{aligned} I(t) &= \langle \hat{i}(t) \rangle = \langle \hat{\mathcal{E}}_p(t)^\dagger \hat{\mathcal{E}}_p(t) \rangle \\ &= (\mathcal{E}_p^+(t) - i\mathcal{E}_p^-(t)) (\mathcal{E}_p^+(t) + i\mathcal{E}_p^-(t)) \\ &= (\mathcal{E}_p^+(t))^2 + (\mathcal{E}_p^-(t))^2. \end{aligned} \quad (2.49)$$

As can be seen, a single photo-detector can only measure the intensity of the pulse. If we are interested in measuring the amplitude or phase quadratures separately we need another technique. For example, if we were to combine our weak probing pulse with a strong classical local oscillator field $|\mathcal{E}_{LO}| \gg |\mathcal{E}_p|$, with a frequency difference $\Delta\omega_{LO}$ and a relative phase ϕ_{LO} on a 50:50 beam-splitter, as shown in Fig. 2.4, we obtain the following

for the two output ports

$$I_1(t) = \frac{1}{2}\mathcal{E}_{LO}^2 + \mathcal{E}_{LO} (\mathcal{E}_p^+(t)\text{Cos}(\Delta\omega_{LO}t + \phi_{LO}) + \mathcal{E}_p^-(t)\text{Sin}(\Delta\omega_{LO}t + \phi_{LO})) \quad (2.50)$$

$$I_2(t) = \frac{1}{2}\mathcal{E}_{LO}^2 - \mathcal{E}_{LO} (\mathcal{E}_p^+(t)\text{Cos}(\Delta\omega_{LO}t + \phi_{LO}) + \mathcal{E}_p^-(t)\text{Sin}(\Delta\omega_{LO}t + \phi_{LO})). \quad (2.51)$$

$$(2.52)$$

If we were to then subtract the two signals we will produce

$$I_{het} = \mathcal{E}_{LO} (\mathcal{E}_p^+(t)\text{Cos}(\Delta\omega_{LO}t + \phi_{LO}) + \mathcal{E}_p^-(t)\text{Sin}(\Delta\omega_{LO}t + \phi_{LO})). \quad (2.53)$$

If $\Delta\omega_{LO} = 0$ then we have homodyne detection, where we obtain the shape of our probe multiplied by the amplitude of the local oscillator and we can view the amplitude or phase quadrature by altering the angle ϕ_{LO} . If $\Delta\omega_{LO} \neq 0$ then this will be equivalent to a constant modulation of our signal which will scan from the amplitude quadrature to the phase quadrature and back, known as heterodyne detection. In this thesis we will mostly work with classical pulses of light and can therefore ignore the noise terms.

2.4 Quantising Atoms

We will not present a derivation of the quantisation of atomic electron orbitals here. For interested readers see, for example, Ref. [204] and a more detailed description of the energy levels of the atoms of interest in this thesis are presented in App. A. Here we will start from the second quantisation of the atomic Hamiltonian, namely where we have quantised atomic levels $|j\rangle$, with corresponding energies $\hbar\omega_j$, to give a Hamiltonian

$$\hat{\mathcal{H}}_a = \sum_j \hbar\omega_j |j\rangle \langle j|, \quad (2.54)$$

where $\langle j|k\rangle = \delta_{jk}$. For simplicity we can define atomic operators

$$\hat{\sigma}_{jk} = |j\rangle \langle k|, \quad (2.55)$$

with terms of the form $\hat{\sigma}_{jj}$ being the atomic populations and $\hat{\sigma}_{jk}$, $j \neq k$ are the coherences. The time-dependence of these operators can be calculated by Eq. 2.11 to be

$$\hat{\sigma}_{jk}(t) = \hat{\sigma}_{jk} \text{Exp}[-i\omega_{jk}t], \quad (2.56)$$

where $\hbar\omega_{jk}$ is the energy difference between the atomic states $|k\rangle$ and $|j\rangle$.

In this thesis we will be performing experiments using ensembles of atoms. In this case we need to define locally-averaged atomic operators by taking a cross-section of the ensemble with length δz , area A , and with an atomic density n to give [212]

$$\hat{\sigma}_{jk}(z, t) = \frac{1}{nA\delta z} \sum_{z_1 \in \delta z} \hat{\sigma}_{jk}^l(z_1, t), \quad (2.57)$$




Figure 2.5: The Three-Level Atom, used throughout this thesis for interaction of a weak probe field $\hat{\mathcal{E}}$ on the $|1\rangle \rightarrow |3\rangle$ transition (dashed line) with coupling strength g , and a strong coupling field with Rabi frequency Ω_c on the $|2\rangle \rightarrow |3\rangle$ transition (solid line). The coupling field is detuned an amount Δ from resonance, and the probe is detuned an amount $\Delta - \delta$, where Δ is known as the one-photon detuning and δ is the two-photon detuning.

where $nA\delta z = N_z$, the atomic population in the slice z . The total number of atoms in an ensemble of length L is therefore given by $N = L N_z$.

2.5 The Interaction of Atoms and Light

We now come to the interaction of atoms with the light field, starting with determining the form of the interaction Hamiltonian.

2.5.1 The Interaction Hamiltonian

We can approximate the interaction Hamiltonian as the potential energy of the atom's electric dipole in the light field to give

$$\hat{\mathcal{V}} = \sum_j e \mathbf{r}_j \cdot \hat{\mathbf{E}}(\mathbf{r}_j, t) \quad (2.58)$$

$$= - \sum_j \left(\sum_k d_{jk} \hat{\sigma}_{jk}(\mathbf{r}_j, t) \right) \cdot \hat{\mathbf{E}}(\mathbf{r}_j, t), \quad (2.59)$$

where $d_{jk} = -e \langle j | \hat{\mathbf{r}} | k \rangle$ is the dipole moment for the $|j\rangle \rightarrow |k\rangle$ transition and equals zero if $j = k$. If we assume that the light field contains only a small range of frequencies detuned from an atomic transition, say $\omega_{12} + \Delta$, and we insert the time-dependence of $\hat{\sigma}$ from Eq. 2.56 with the definition of the electric field in terms of $\hat{\mathcal{E}}$ from Eq. 2.36 we will find for a single atom that

$$\begin{aligned} \hat{\mathcal{V}} = & -i\hbar \sqrt{\frac{\omega_{12}}{2\hbar\epsilon_0 V}} \left(-d_{12} \hat{\sigma}_{12} \hat{\mathcal{E}}^\dagger \text{Exp}[i\Delta t] + d_{21} \hat{\sigma}_{21} \hat{\mathcal{E}} \text{Exp}[-i\Delta t] \right. \\ & \left. -d_{12} \hat{\sigma}_{12} \hat{\mathcal{E}} \text{Exp}[-2i\omega_{12}t - i\Delta t] + d_{21} \hat{\sigma}_{21} \hat{\mathcal{E}}^\dagger \text{Exp}[2i\omega_{12}t + i\Delta t] \right) \end{aligned}$$




Figure 2.6: Dressed States. (a) Bare and (b) dressed state picture in the presence of a strong field with Rabi frequency $\Omega_n = g\sqrt{n+1}$, also shows Mollow triplet formation with different colours indicating different frequencies.

$$= -g_{12}\hbar \left(\hat{\sigma}_{12}\hat{\mathcal{E}}^\dagger \text{Exp}[i\Delta t] + \hat{\sigma}_{21}\hat{\mathcal{E}} \text{Exp}[-i\Delta t] \right), \quad (2.60)$$

where we have performed the rotating-wave approximation to remove the fast oscillating terms ($2i\omega_{12}t$) and rotated into a frame where $d_{12} = i|d_{12}|$ to produce the Jaynes-Cummings interaction Hamiltonian, with a atom-light coupling strength

$$g_{12} = \sqrt{\frac{\omega_{12}}{2\hbar\epsilon_0 V}} |d_{12}|. \quad (2.61)$$

If we were to have N two-level atoms in a length L in the z -direction (i.e., the direction of propagation of the light field) then, in the continuum limit, we will have

$$\hat{\mathcal{V}} = -\frac{\hbar g N}{L} \int_0^L dz \left(\hat{\sigma}_{21}(z,t)\hat{\mathcal{E}}(z,t)\text{Exp}[-i\Delta t] + \hat{\sigma}_{12}(z,t)\hat{\mathcal{E}}^\dagger(z,t)\text{Exp}[i\Delta t] \right). \quad (2.62)$$

Similarly, if we have two light fields, $\hat{\mathcal{E}}_{23}$ interacting with the $|2\rangle \rightarrow |3\rangle$ transition, detuned from resonance by Δ_{23} , and another field $\hat{\mathcal{E}}_{13}$ interacting with the $|1\rangle \rightarrow |3\rangle$ transition, detuned from resonance by Δ_{13} , then we can describe the situation with the following Hamiltonian

$$\begin{aligned} \hat{\mathcal{V}} = & -\frac{\hbar N}{L} \int_0^L dz \left(\hat{\sigma}_{32}(z,t)\Omega_c \text{Exp}[-i\Delta_{23}t] \right. \\ & \left. + g\hat{\sigma}_{31}(z,t)\hat{\mathcal{E}}_{13} \text{Exp}[-i\Delta_{13}t] + \text{H.C.} \right), \end{aligned} \quad (2.63)$$

where H.C. is the Hermitian conjugate. This is illustrated in Fig. 2.5 where we define $\Omega_c = g_{23}\langle\hat{\mathcal{E}}_{23}\rangle$ as a classical Rabi frequency that be treated as a complex number (as its phase is not important), label $\Delta = \Delta_{23}$ as the one-photon detuning and $\delta = \Delta_{23} - \Delta_{13}$ as the two-photon detuning. We also assume that state $|3\rangle$ is the excited state, with the two ground-states $|1\rangle$ and $|2\rangle$ being much closer in energy, for reasons that will become clear later.

2.5.2 The Dressed State Picture

If we revert for a minute to a single two-level atom with a single-mode light field then we can write the total system Hamiltonian as

$$\begin{aligned}\hat{\mathcal{H}} &= \hat{\mathcal{H}}_l + \hat{\mathcal{H}}_a + \hat{\mathcal{V}} \\ &= \hbar\omega_{12}\hat{\sigma}_{22} + \hbar\omega_l\hat{a}^\dagger\hat{a} + g\hat{a}^\dagger\hat{\sigma}_{21} + g\hat{a}\hat{\sigma}_{12},\end{aligned}\quad (2.64)$$

where ω_l is the frequency of the light field. If we assume $\omega_l = \omega_{12}$ then we see that the above Hamiltonian will couple states of the form $|n\rangle_l|1\rangle_a$ and $|n+1\rangle_l|2\rangle_a$ via

$$\langle n|_l\langle 1|_a\hat{\mathcal{V}}|n+1\rangle_l|2\rangle_a = g\sqrt{n+1},\quad (2.65)$$

where the subscripts l and a refer to the light and atomic states respectively, both states having the same energy. We can therefore write the energy eigenstates of the free Hamiltonian (i.e., $\hat{\mathcal{H}}_l + \hat{\mathcal{H}}_a$) as

$$|+_n\rangle = \frac{1}{\sqrt{2}}(|n\rangle_l|1\rangle_a + |n+1\rangle_l|2\rangle_a)\quad (2.66)$$

$$|-_n\rangle = \frac{1}{\sqrt{2}}(|n\rangle_l|1\rangle_a - |n+1\rangle_l|2\rangle_a),\quad (2.67)$$

with relative energies $E_\pm = \pm\hbar g\sqrt{n+1}$ from the initial state energy (i.e., $\hbar\omega_{12}$ for the excited state). This is known as the dressed-state picture and this level-splitting effect responsible for such phenomena as Mollow triplets [213]. This is illustrated in Fig. 2.6. It can also be explained in terms of the ac Stark effect (see Ch. 10).

2.5.3 Time-Evolution

Now that we have the form of the interaction Hamiltonian, we can determine the time-evolution of the atomic and electric field operators. Using the fact that the individual and continuum atomic operators have the commutation relation

$$[\hat{\sigma}_{jk}(z), \hat{\sigma}_{kl}(z')] = \frac{1}{N_z}\hat{\sigma}_{jl}(z)\delta(z-z')\quad (2.68)$$

$$[\hat{\sigma}_{jk}(z), \hat{\sigma}_{kj}(z')] = \frac{1}{N_z}(\hat{\sigma}_{jj}(z) - \hat{\sigma}_{kk}(z))\delta(z-z'),\quad (2.69)$$

we can find the expectation values of the atomic operators using the Heisenberg-Langevin equations from Eq. 2.12 using $\hat{\mathcal{V}}$ to give

$$\frac{\partial}{\partial t}\hat{\sigma}_{11} = \gamma\hat{\sigma}_{33} + \gamma_c(\hat{\sigma}_{22} - \hat{\sigma}_{11}) + ig\left(\hat{\mathcal{E}}^\dagger\hat{\sigma}_{13} - \hat{\mathcal{E}}\hat{\sigma}_{31}\right) + \hat{F}_{11}\quad (2.70)$$

$$\frac{\partial}{\partial t}\hat{\sigma}_{22} = \gamma\hat{\sigma}_{33} + \gamma_c(\hat{\sigma}_{11} - \hat{\sigma}_{22}) + i\left(\Omega_c^*\hat{\sigma}_{23} - \Omega_c\hat{\sigma}_{32}\right) + \hat{F}_{22}\quad (2.71)$$

$$\frac{\partial}{\partial t}\hat{\sigma}_{33} = -2\gamma\hat{\sigma}_{33} + i\left(g\hat{\mathcal{E}}\hat{\sigma}_{31} + \Omega_c\hat{\sigma}_{32} - \text{H.C.}\right) + \hat{F}_{33}\quad (2.72)$$

$$\frac{\partial}{\partial t}\hat{\sigma}_{12} = -(\gamma_o + \gamma_c + i\delta)\hat{\sigma}_{12} + i\left(\Omega_c^*\hat{\sigma}_{13} - g\hat{\mathcal{E}}\hat{\sigma}_{32}\right) + \hat{F}_{12}\quad (2.73)$$

$$\frac{\partial}{\partial t}\hat{\sigma}_{13} = -(\gamma + \gamma_o/2 + \gamma_c/2 + i\Delta)\hat{\sigma}_{13} + ig\hat{\mathcal{E}}(\hat{\sigma}_{11} - \hat{\sigma}_{33}) + i\Omega_c\hat{\sigma}_{12} + \hat{F}_{13}\quad (2.74)$$

$$\frac{\partial}{\partial t} \hat{\sigma}_{32} = -(\gamma + \gamma_o/2 + \gamma_c/2 + i(\Delta - \delta)) \hat{\sigma}_{32} + i\Omega_c^* (\hat{\sigma}_{33} - \hat{\sigma}_{22}) - ig\hat{\mathcal{E}}^\dagger \hat{\sigma}_{12} + \hat{F}_{23} \quad (2.75)$$

where we have dropped the explicit time- and spatial-dependence for ease of viewing. The γ , γ_o , and γ_c terms denote, respectively, the decay rate from $|3\rangle$ to the other two states (taken to be the same for the $|3\rangle \rightarrow |1\rangle$ and $|3\rangle \rightarrow |2\rangle$ transitions), the ground-state decoherence rate, and ground-state population exchange (through collisions etc.). These are illustrated in Fig. 2.5. The decay rate depends on the atom-light coupling strength g , and is much larger than the two ground-state decoherence terms $\gamma \gg \gamma_o, \gamma_c$. Also, γ_o is normally much larger than γ_c [139] and we will therefore ignore γ_c in future analysis. The Langevin operators \hat{F}_{jk} are required to preserve the commutation relations for the system in the presence of spontaneous emission-induced damping. Here we will ignore these terms, as they solely add vacuum noise in the regimes of interest to us [214, 137].

We can determine the dynamics of $\hat{\mathcal{E}}_{13}$ using the Heisenberg equation of motion (Eq. 2.11) with the relevant Hamiltonian being the sum of the interaction Hamiltonian (Eq. 2.63) and the multi-mode light field Hamiltonian (Eq. 2.40) to give

$$\begin{aligned} \frac{d}{dt} \hat{\mathcal{E}}_{13} &= \frac{-i}{\hbar} [\hat{\mathcal{E}}_{13}, \hat{\mathcal{H}}_l + \hat{\mathcal{V}}] \\ &= \frac{-i}{\hbar} \left(i \int dk' \left(\sum_k \hbar\omega_k \text{Exp}[i\mathbf{k}' \cdot \mathbf{r}] [\hat{a}_{k'}, \hat{a}_k^\dagger \hat{a}_k] \right) \right. \\ &\quad \left. - \frac{N}{L} \int dz (g\hat{\sigma}_{13} [\hat{\mathcal{E}}_{13}, \hat{\mathcal{E}}_{13}^\dagger]) \right), \end{aligned} \quad (2.76)$$

where we have kept only the terms whose commutator is non-zero and used the definition of $\hat{\mathcal{E}}$ from Eq. 2.35. Now if we use the commutator relations for $\hat{\mathcal{E}}$ (Eq. 2.37) and \hat{a}_k (Eq. 2.15) we obtain

$$\frac{d}{dt} \hat{\mathcal{E}}_{13} = -\omega_{13} \hat{\mathcal{E}}_{13} + igN \hat{\sigma}_{13}. \quad (2.77)$$

We can also determine the z derivative of $\hat{\mathcal{E}}_{13}$ as follows

$$\begin{aligned} \frac{\partial}{\partial z} \hat{\mathcal{E}}_{13} &= i \int dz \left(\hat{a}_k \frac{\partial}{\partial z} \text{Exp}[i\mathbf{k} \cdot \mathbf{r}] \right) \\ &= (ik_z) \hat{\mathcal{E}}_{13}. \end{aligned} \quad (2.78)$$

We can now combine these two equations to give us what is known as the Maxwell equation for $\hat{\mathcal{E}}_{13}$

$$\begin{aligned} \frac{d}{dt} \hat{\mathcal{E}}_{13} + c \frac{\partial}{\partial z} \hat{\mathcal{E}}_{13} &= igN \hat{\sigma}_{13} + i(\omega_{13} - ck_z) \hat{\mathcal{E}}_{13} \\ &= igN \hat{\sigma}_{13} + \frac{ic}{2k_{13}} \nabla_{x,y}^2 \hat{\mathcal{E}}_{13}, \end{aligned} \quad (2.79)$$

where we have made the paraxial approximation $k_z \gg k_x, k_y$ and therefore that $k_{13} \approx k_z$. We can then approximate

$$\omega_{13} - ck_z = c \left(k_{13} - \frac{k_z}{\sqrt{k_x^2 + k_y^2 + k_z^2}} \right)$$




Figure 2.7: Bloch Sphere. (a) Bloch sphere showing the interaction between a two-level atom and a light field (i) on resonance (blue) and (ii) off resonance (red). (b) Bloch vectors showing the effects of (i) population decay with time constant T_1 (blue) and (ii) decoherence with time constant T_2^* (red). (c) Absorption spectrum of an atomic ensemble with individual, homogeneous, linewidths γ and overall, inhomogeneous, linewidth $\delta\omega$.

$$\approx \frac{1}{2k_{13}} (k_x^2 + k_y^2).$$

For most of this thesis we will make a plane-wave approximation, i.e., $\text{Exp}[i\mathbf{k} \cdot \mathbf{r}] = \text{Exp}[ik_z z]$, and therefore the second term from Eq. 2.79 will vanish, leaving us with

$$\left(\frac{d}{dt} + c \frac{\partial}{\partial z} \right) \hat{\mathcal{E}}(z, t) = igN \hat{\sigma}_{13}(z, t). \quad (2.80)$$

We will use this equation to help determine the time-evolution of single-mode storage in both EIT and GEM in the next chapter. However, the $\nabla_{x,y}^2 \hat{\mathcal{E}}_{13}$ term from Eq. 2.79 will be important for TEM modes greater than TEM₀₀, and therefore image storage (see Ch. 8).

2.5.4 Optical Bloch Equations

The optical Bloch equations and corresponding Bloch sphere visualisations, already seen in the previous chapter, are an elegant way of describing the interaction of a light field with two-level atoms. We start by defining a Bloch vector $\rho = (u, v, w)$, where the components are described in terms of the atomic populations and coherences by [204]

$$u = \sigma_{12} + \sigma_{21} \quad (2.81)$$

$$v = i(\sigma_{21} - \sigma_{12}) \quad (2.82)$$

$$w = \sigma_{11} - \sigma_{22}, \quad (2.83)$$

where we have removed the hats for ease of notation. If we have a light field $\mathcal{E}(t)$ incident on our two-level atom, detuned an amount Δ from resonance, the evolution will be described by its interaction with the torque vector $\mathbf{\Omega} = (-g\mathcal{E}(t), 0, -\Delta)$ in a frame rotating at the atomic transition frequency ω_{12} by

$$\frac{d}{dt} \rho = \mathbf{\Omega} \times \rho. \quad (2.84)$$

This gives the following equations for the time-evolution of the individual components

$$\frac{d}{dt}u = -\Delta v \quad (2.85)$$

$$\frac{d}{dt}v = \Delta u - g\mathcal{E}(t)w \quad (2.86)$$

$$\frac{d}{dt}w = g\mathcal{E}(t)v. \quad (2.87)$$

From this set of equations it can be seen that if the atom is subjected to a constant light field, on resonance, then it will simply oscillate between the ground and excited states at the Rabi frequency $\Omega_c = g\mathcal{E}$. This type of behaviour is called Rabi flopping and is illustrated in Fig. 2.7(a)(i). This assumes no spontaneous emission or decoherence. If, on the other hand, we were to excite the atom with an on-resonance pulse of area $\mathcal{A}(t) = \int_{-\infty}^t dt' g\mathcal{E}(t')$ then, after time t , we will find the system in the state

$$u(t) = 0 \quad (2.88)$$

$$v(t) = \text{Sin}[\mathcal{A}(t)] \quad (2.89)$$

$$w(t) = \text{Cos}[\mathcal{A}(t)]. \quad (2.90)$$

$$(2.91)$$

It follows from these equations that, if $\mathcal{A}(t) = \pi/2$ then $\rho(t) = (0, 1, 0)$ and, if $\mathcal{A}(t) = \pi$ then $\rho(t) = (0, 0, 1)$. These are the $\pi/2$ - and π -pulses mentioned in the introduction.

If the light is not on resonance (i.e., $\Delta \neq 0$) we can perform a coordinate change to help with determining the system dynamics by rotating by an angle $\phi = \text{Tan}^{-1}[\Delta/g\mathcal{E}(t)]$ about the v axis. The time-evolution in these new coordinates is given by

$$\frac{d}{dt}u' = 0 \quad (2.92)$$

$$\frac{d}{dt}v' = -\Omega_c(t, \Delta)w \quad (2.93)$$

$$\frac{d}{dt}w' = \Omega_c(t, \Delta)v., \quad (2.94)$$

where $\Omega_c(t, \Delta) = \sqrt{g^2\mathcal{E}^2(t) + \Delta^2}$ is the off-resonance Rabi frequency. The evolution of ρ for a constant detuned field is shown in Fig. 2.7(a)(ii).

We will now ask the question: what will happen to the Bloch vector dynamics once the light field is removed? Initially assuming that the excitation pulse was on resonance, the Bloch vector will remain stationary according to the above equations. This is not the case in practice, however, due to spontaneous emission with rate γ and decoherence with rate γ_o . As stated in the previous section, γ is responsible for population decay from $|2\rangle \rightarrow |1\rangle$ (also called longitudinal decay) with a time-constant T_1 . γ_o , meanwhile, is responsible for dephasing (transverse decay, due to scattering from the surface of the sphere) with time-constant T_2^* . We can phenomenologically add these terms to the Bloch evolution equation to produce [203]

$$\frac{d}{dt}u = -\Delta v - \frac{u}{T} \quad (2.95)$$

$$\frac{d}{dt}v = \Delta u - g\mathcal{E}(t)w - \frac{v}{T} \quad (2.96)$$

$$\frac{d}{dt}w = g\mathcal{E}(t)v - wT_1, \quad (2.97)$$

where $T = T_1 + T_2^*$. The T_1 and T_2^* processes are illustrated in Fig. 2.7(b). As can be seen, decoherence does not lead to a loss of population, but leads to a reduction in the length of the Bloch vector, while spontaneous emission leads to population decay as well as a change in length of the Bloch vector. The reduction of the length of the Bloch vector is significant, as coherent superpositions states sit on the surface of the sphere while statistical mixtures sit inside the sphere.

If, for a moment, we neglect the damping terms and assume the atom was illuminated by off-resonance $\pi/2$ -pulse of length τ , short compared to the $1/\Delta$, we find that²

$$\rho(t) = (\text{Sin} [\Delta t], \text{Cos} [\Delta t], 0). \quad (2.98)$$

That is, the Bloch vector rotates around the $u - v$ plane at a rate determined by the detuning of the light.

Now let us consider an ensemble of two-level atoms. The absorption profile of an individual atom is Lorentzian with a full-width-half-maximum (FWHM) at its lowest limit of γ , i.e.,

$$L(\omega) = \frac{\gamma}{2\pi} \frac{1}{(\omega - \omega_{12})^2 + (\gamma/2)^2}. \quad (2.99)$$

This is known as the homogeneous broadening of the atom, as the width will be the same for all atoms. However, the centre absorption frequency will vary from atom to atom due to, for example, Doppler shifts or inhomogeneous electric or magnetic fields. This leads to an absorption bandwidth for an ensemble of atoms that is greater than the bandwidth of an individual atom. This is called inhomogeneous broadening. The difference between homogeneous and inhomogeneous broadening is illustrated in Fig. 2.7(c). Inhomogeneous broadening is desirable for the absorption of light, as we would like to absorb as large a range of frequencies as possible. However, if we consider inhomogeneous broadening in terms of Eq. 2.98, where each individual dipole will rotate at a slightly different frequency due to the inhomogeneous broadening, then we see that the sum of all the Bloch vectors, known as the macroscopic coherence, will destructively interfere and head to zero. We can investigate the macroscopic coherence mathematically by looking at the ensemble polarisation

$$\mathcal{P}(t) \propto \int d\Delta' \text{Re} \{ (u(t, \Delta') + iv(t, \Delta')) \text{Exp} [i\omega_{12}] \} I(\Delta'), \quad (2.100)$$

where $I(\Delta)$ is the inhomogeneous broadening function. If we were to assume a Lorentzian inhomogeneous broadening with FWHM $\delta\omega$ and centre frequency $\omega_{12} + \Delta$ then the above equation simplifies to

$$\mathcal{P}(t) = \mathcal{P}(0) \text{Cos} [(\omega_{12} + \Delta)t] \text{Exp} [-(\delta\omega)t]. \quad (2.101)$$

As can be seen, there are two components to the polarisation: a cosine oscillating close to the atomic transition frequency and a slow decay, determined by the inhomogeneous linewidth. If we could find a way to cause the individual Bloch vectors to realign then we could regain the macroscopic coherence and, by a reversal of the absorption process, cause the coherent emission of a pulse from the ensemble. This is the basis of the photon

²A $\pi/2$ -pulse is not necessary to observe this behaviour, but is used for simplicity.

echo techniques discussed in the next chapter.

2.6 Conclusions

In this chapter we have presented the basic theory needed to understand the ideas presented in the previous chapter with regards to quantum phenomena including coherence, entanglement and teleportation, as well as providing an introduction to atom-light interaction theory as a spring-board for the development of quantum memory theory presented in the next chapter.

Quantum Memory Theory

Sam: "Oh, boy ... " - *Quantum Leap* (repeated line), NBC, 1989-1993, TV Series.

3.1 A Memorable Introduction

In this chapter we will use the basic theory presented in the previous chapter to construct the workings of the two quantum memory protocols used in this thesis: Electromagnetically-induced transparency (EIT) and photon-echo-based memories. These are presented in Sec. 3.3 and Sec. 3.4 respectively. But first we will take a more in-depth look at a few of the quantum memory metrics discussed in the introduction.

3.2 Quantum Memory Metrics Take II

Here we will give a more solid mathematical basis for some of the quantum memory performance metrics presented in Sec. 1.4.

3.2.1 Efficiency, Fidelity and T-V Diagrams

Let us start by defining efficiency as the ratio of the number of photons in the output pulse compared to the input, i.e.,

$$\varepsilon_m = \frac{\langle \hat{\mathcal{E}}_{out}^\dagger \hat{\mathcal{E}}_{out} \rangle}{\langle \hat{\mathcal{E}}_{in}^\dagger \hat{\mathcal{E}}_{in} \rangle}. \quad (3.1)$$

Classically, a non-unit efficiency will simply lead to a loss of signal, which could be recovered by amplification. However, from a quantum-mechanical picture, loss in a system is equivalent to placing a beam-splitter in the optical path with transmissivity $\sqrt{\varepsilon_m}$. This is




Figure 3.1: Loss from Quantum Systems, equivalent to a beam-splitter (BS) coupling in vacuum noise X_{vac} . For quantum memories the transmissivity of the BS will be given by $\sqrt{\varepsilon_m}$, where ε_m is the memory efficiency.

an issue for the recalled quantum state as a quanta of vacuum noise X_{vac} will be introduced through the other entry port of the beam-splitter, degrading the output state. This is illustrated in Fig. 3.1. We also cannot deterministically amplify an arbitrary state without increasing the noise further [215]. In a quantum optics experiment losses may come from absorption or scatter off optics, or from inefficient detection (both the intrinsic detector efficiency and visibility for heterodyne/homodyne measurements). Except when performing noise measurements, in this thesis we will concern ourselves only with the memory efficiency. However, if we were to investigate the storage of, for example, squeezed states then these other sources of loss would be an issue.

The unconditional fidelity of a state is linked to the efficiency. This was illustrated in Fig. 1.7 for coherent states. Mathematically we can describe the fidelity of a memory by the overlap of the input and output states, i.e.,

$$\mathcal{F} = |\langle \psi_{out} | \psi_{in} \rangle|^2. \quad (3.2)$$

Fidelity is a state-dependent measure. For instance, for a coherent state with mean photon number $|\alpha|^2 = n_p$, the fidelity of the output state as a function of the memory efficiency, assuming a linear memory process, will be [216]

$$\mathcal{F}_{n_p}(\epsilon_m) = \frac{1}{1 + n_p (1 - \sqrt{\epsilon_m})^2}. \quad (3.3)$$

From this equation it can be seen that, even for different sized coherent states, the fidelity of a fixed-efficiency memory will be different, decreasing as n_p increases. Even so, much theoretical work has been done investigating bounds with fidelity. The classical fidelity bound - the maximum fidelity for a classical measure-and-reproduce memory - for coherent states is given by [217, 218]

$$\mathcal{F}_c(n_p) \leq \frac{1 + n_p}{2n_p + 1}. \quad (3.4)$$

For large mean photon numbers this tends to 50%, the classical limit discussed in the introduction. To break the no-cloning limit for coherent states we need a fidelity of greater than 68% [97]. Using Eq. 3.3 we can see that for a coherent state with $n_p = 10$ we require an efficiency of 60% and for $n_p = 100$ we require 87%.

T-V diagrams [219] were proposed as a state-independent method of measuring how well an output state resembles an input state, originally for quantum non-demolition measurements [98] followed by quantum teleportation [99]. We start here by assuming that a light field has a coherent information term and a much smaller fluctuating noise term, i.e.,

$$X_{in/out}^\pm = \mathcal{E}_{in/out}^\pm + \delta \hat{X}_{in/out}^\pm, \quad (3.5)$$

where the \mathcal{E}_{in}^\pm etc. terms refer to the amplitude or phase quadrature of the light state. We can then define the variance of these fields V^\pm using Eq. 2.9. Once we have these terms, we can define the signal to noise ratio of the state

$$\mathcal{R}_{in/out} = \frac{4 \left(\mathcal{E}_{in/out}^\pm \right)^2}{V_{in/out}^\pm}. \quad (3.6)$$

Next we define the conditional variances and signal transfer coefficients to be, respectively,

$$V_c^\pm = V_{out}^\pm - \frac{|\langle \hat{X}_{in}^\pm \hat{X}_{out}^\pm \rangle|^2}{V_{in}^\pm} \quad (3.7)$$

$$T^\pm = \frac{\mathcal{R}_{out}^\pm}{\mathcal{R}_{in}^\pm}. \quad (3.8)$$

Finally, we define $V = \sqrt{V_c^+ V_c^-}$ and $T = T^+ + T^-$. To produce the best possible copy of the input state (i.e., to be in the no-cloning regime) we need to have high efficiency without adding any noise. On a T-V diagram this region is represented by $V < 1$ (low noise) and $T > 1$ (high efficiency) [96], while a classical measure-and-reproduce memory cannot beat either of these bounds [99].

3.2.2 Bandwidths and Coherence Times

We will now have a very quick look into mathematically describing the bandwidth and coherence times of a memory. In this thesis we will mostly work with Gaussian pulses, i.e., pulses with a temporal profile

$$\mathcal{E}(t) = \mathcal{E} \text{Exp} \left[\frac{-t^2}{2\sigma_t^2} \right], \quad (3.9)$$

where σ_t is the Gaussian temporal standard deviation. The full-width-half-maximum (FWHM) of such a pulse is given by $2\sigma_t\sqrt{2\ln[2]}$. By performing a Fourier transform of such a pulse we can determine its bandwidth. As the Fourier transform of a Gaussian is another Gaussian, with a new standard deviation $\sigma_\omega = 1/2\pi\sigma_t$, we normally define the pulse's bandwidth \mathcal{B}_p as its FWHM in frequency, i.e.,

$$\mathcal{B}_p = \frac{\sqrt{2\ln[2]}}{\pi\sigma_t}. \quad (3.10)$$

Of course, in practise we want to store all frequency components of a pulse. To store 99% of the pulse would require a memory bandwidth of $6\sigma_\omega$.

We found in Sec. 2.5.4 that a state will decay exponentially due to spontaneous emission and decoherence, i.e., $\text{Exp}[-t/T]$, where $T = T_1 + T_2^*$. Other effects, such as diffusion of atoms out of the interaction volume defined by the light field, are expected to follow a Gaussian decay, i.e., $\text{Exp}[-t^2/T^2]$. The total decay will be a combination of the two, i.e., $\text{Exp}[-t/T_e] \cdot \text{Exp}[-t^2/T_g^2]$. Normally one decay term is much greater than the other, for instance $T_e \gg T_g$, and so the decay of a state can be described with just one of these terms. We then define the coherence time of the memory by this limiting time. We can also define the delay-bandwidth product (DBP) using these terms as $\mathcal{B}_p T$. Alternative versions of the delay-bandwidth product simply use the ratio of maximum storage time t_s and pulse length t_p , or place limits on the efficiency for t_s (i.e., $\epsilon_m(t < t_s) > 50\%$).

3.3 Electromagnetically-Induced Transparency Theory

We now return to studying the mathematical behaviour of quantum memories, firstly looking at electromagnetically-induced transparency (EIT).

3.3.1 Slowing Light with EIT

The level scheme for EIT was presented in Fig. 1.8(a). It requires two ground-states $|1\rangle$ and $|2\rangle$, an excited-state $|3\rangle$ and two optical fields, the weak probe $\hat{\mathcal{E}}_p$ on the $|1\rangle \rightarrow |3\rangle$ transition and a coupling field with Rabi frequency Ω_c on the $|2\rangle \rightarrow |3\rangle$ transition. Firstly, it should be noted that EIT is a coherent population trapping effect (for instance, see Ref. [134]), caused by destructive Fano interference between the light and atomic fields [220], and not due to dressed-state splitting caused by the coupling field (see Sec. 2.5.2). The latter effect results in what is known as Autler-Townes splitting [221, 222] and, while both lead to enhanced transmission for the probe, they should be considered separately¹

To analyse the situation mathematically, let us start from the time-evolution equations for the atomic operators and probe field (Eq. 2.70-2.75 and Eq. 2.80 respectively). If we start our analysis by assuming the system is in a steady state (i.e., $\partial\hat{\sigma}/\partial t = 0$) in a reference frame moving at the speed of light ($t' = t - z/c$), and Fourier transforming our operators using

$$\hat{O}(\omega) = \frac{1}{2\pi} \int_0^\infty dt \hat{O}(t) \text{Exp}[-i\omega t], \quad (3.11)$$

we obtain the following equation for the spatial evolution of $\hat{\mathcal{E}}_p$

$$\frac{\partial}{\partial z} \hat{\mathcal{E}}_p(z, \omega) = -\mathcal{T}(\omega) \hat{\mathcal{E}}_p(z, \omega) \quad (3.12)$$

giving a solution

$$\hat{\mathcal{E}}_p(L, \omega) = \hat{\mathcal{E}}_p(0, \omega) \text{Exp}[-\mathcal{T}(\omega)L]. \quad (3.13)$$

$\mathcal{T}(\omega)$ is the transfer function for the medium. The susceptibility of the EIT medium (i.e., where $\Delta = 0$) is given by $c\mathcal{T}/igN$ to be [223]

$$\chi(\omega) = g \left[\frac{4\delta \left(|\Omega_c|^2 - 4\delta^2 \right) - 4\delta\gamma_o^2 + i \left(8\delta^2 (\gamma + \gamma_o) + 2\gamma_o \left(|\Omega_c|^2 + \gamma_o\gamma \right) \right)}{\left| |\Omega_c|^2 + (\gamma + \gamma_o + 2i\delta) (\gamma_o + 2i\delta) \right|^2} \right], \quad (3.14)$$

assuming that we are in the weak probe regime where $\langle \hat{\sigma}_{11} \rangle \approx 1$ and $g^2 \langle \hat{\mathcal{E}} \rangle^2 \approx 0$. The susceptibility is very important for EIT, as it tells us a lot about the dynamics of the system. For instance, the imaginary component determines the absorption of the system, via Eq. 3.13, and the real component determines the refractive index, via

$$\mathcal{R}(\omega) = 1 + \frac{1}{2} \text{Re} [\chi(\omega)], \quad (3.15)$$

and therefore the dispersion. These are shown in Fig. 1.8(b) and (c) respectively. As can be seen, when there is a coupling field present we have enhanced transmission of the probe, accompanied by a sharp change in refractive index, indicating a large drop in probe group

¹For a more detailed discussion of the differences between the two, see Sec. 11.3.3.

velocity v_g . To determine the group velocity mathematically we again make the adiabatic and weak-probe approximations, as well as ignore the decay terms for simplicity and assuming $\Delta = \delta = 0$, to obtain

$$\hat{\sigma}_{13} = \frac{-i}{\Omega_c(t)} \frac{\partial}{\partial t} \hat{\sigma}_{12} \quad (3.16)$$

$$\hat{\sigma}_{12} = \frac{g\hat{\mathcal{E}}(z,t)}{\Omega_c(t)} \quad (3.17)$$

$$\left(\frac{d}{dt} + c \frac{\partial}{\partial z} \right) \hat{\mathcal{E}}(z,t) = \frac{-g^2 N}{\Omega_c(t)} \frac{\partial}{\partial t} \left(\frac{\hat{\mathcal{E}}(z,t)}{\Omega_c(t)} \right). \quad (3.18)$$

If we maintain the assumption that the coupling field is constant, then Eq. 3.18 becomes

$$\left(\frac{d}{dt} + \frac{c\Omega_c^2}{1+g^2N} \frac{\partial}{\partial z} \right) \hat{\mathcal{E}}(z,t) = 0, \quad (3.19)$$

and therefore the group velocity will be

$$v_p = \frac{c\Omega_c^2}{1+g^2N} \approx \frac{c\Omega_c^2}{g^2N} = \frac{\Omega_c^2 L}{\gamma \text{OD}}, \quad (3.20)$$

where we have assumed that $g^2 N / \Omega_c^2 \ll 1$, with L being the length of the memory, and defining the optical depth (OD) of the ensemble to be

$$\text{OD} = \frac{g^2 N L}{c\gamma}. \quad (3.21)$$

We can also determine the transmission bandwidth of the EIT memory from the linewidth of the transparency feature shown in Fig. 1.8(b) to be

$$\mathcal{B}_{EIT} = \gamma_o + \frac{\Omega_c^2}{\gamma\sqrt{\text{OD}}}. \quad (3.22)$$

From these equations we can see that as we decrease the Rabi frequency (proportional to the intensity of the coupling field) we both reduce the group velocity, allowing longer pulses to fully fit within the memory length, and reduce the frequency range that will fit inside the memory. This was the issue mentioned in the introduction with regards to storing multiple pulses using EIT: to store many pulses we need the total temporal length $\tau \propto 1/v_p$ to physically fit inside the memory while the bandwidth would have to scale as n/τ , where n is the number of temporal modes in the pulse train. Another way of viewing this is that the product $\mathcal{B}_{EIT}/v_p \propto \sqrt{\text{OD}}$ and therefore large ODs are needed to be able to store many modes.

3.3.2 EIT Polariton and Storing Light

We now discuss the mathematics behind the storage of light using EIT, as this requires a dynamic $\Omega_c(t)$. From the above information one can imagine that if, after a pulse has entered the storage medium with enhanced transparency, and is travelling at a group velocity $v_p \propto \Omega_c^2(t)$ then, if we were to reduce $\Omega_c(t) \rightarrow 0$, the pulse would cease to move and therefore be stored in the memory. If we return to Eq. 3.18, we can solve this without

removing the time dependence of $\Omega_c(t)$ by introducing the states [120]

$$\hat{\Psi}_{\pm}(z, t) = \text{Cos}[\theta(t)] \hat{\mathcal{E}}(z, t) \pm \text{Sin}[\theta(t)] \sqrt{N} \hat{\sigma}_{12}(z, t) \quad (3.23)$$

$$\text{Cos}[\theta(t)] = \frac{\Omega_c(t)}{\sqrt{\Omega_c^2(t) + g^2 N}} \quad (3.24)$$

$$\text{Sin}[\theta(t)] = \frac{gN}{\sqrt{\Omega_c^2(t) + g^2 N}}. \quad (3.25)$$

We call these states polaritons (bosonic quasi-particles) and they obey

$$\left[\frac{d}{dt} + c \text{Cos}^2[\theta(t)] \frac{\partial}{\partial z} \right] \hat{\Psi}_-(z, t) = 0 \quad (3.26)$$

$$\hat{\Psi}_+(z, t) = 0. \quad (3.27)$$

The first of these is called the dark-state polariton, the second is called the bright-state polariton and is a trivial solution. Again, it can be seen that the group velocity $v_g = c \text{Cos}^2[\theta(t)]$. As $\Omega_c(t) \rightarrow 0$ we see firstly that $v_p \rightarrow 0$, along with the cosine term, while the sine term tends to one. The cosine and sine represent the light and atomic components of the polariton respectively and, therefore, reducing the coupling power causes information to be removed from the light field and stored in the atomic component of the polariton. Similarly, increasing the coupling power from zero leads to mapping of the information from the atoms to the light field. Polaritons are an elegant and useful way of representing a light-atom system, as we will see again when discussing GEM.

3.4 Gradient Echo Memory Theory

Here we present the theory of light storage using the gradient echo memory (GEM) scheme, starting by focusing on two-level systems.

3.4.1 Echo Memories

The goal in all echo-memory-style technique is to recreate the initial macroscopic coherence of dipoles, discussed in Sec. 2.5.4, to time-reverse the absorption process and therefore create an echo of the stored information. The workings of this type of memory stem from the atomic polarisation (Eq. 2.100), which can be rewritten and related to the Maxwell equation for $\hat{\mathcal{E}}$ as follows

$$\left(\frac{d}{dt} + c \frac{\partial}{\partial z} \right) \hat{\mathcal{E}}_{13}(z, t) = igN \int d\Delta I(\Delta) \hat{\sigma}_{13}(z, t, \Delta), \quad (3.28)$$

where $I(\Delta)$ is the inhomogeneous broadening function. As was discussed in Sec. 2.5.4, the inhomogeneous broadening will lead to a loss of macroscopic coherence and therefore the integral will head to zero.

There are a few ways of realigning the dipoles and we will quickly run through these here. Firstly, the Hahn photon echo scheme. Here the ensemble is excited by a $\pi/2$ -pulse. From Eq.s 2.88-2.90 we see that after the absorption the Bloch vector will be $\rho(0) = (0, 1, 0)$. Then, the different detunings will lead to rotation around the Bloch

sphere meaning that after time τ the Bloch vector will be

$$\rho(\tau, \Delta) = (\text{Sin}[\Delta\tau], \text{Cos}[\Delta\tau], 0), \quad (3.29)$$

(ignoring decoherence effects). Now, if we were to apply a π -pulse $(u, v, 0) \rightarrow (u, -v, 0)$. Evolution after this pulse will be described by

$$\rho(t > \tau, \Delta) = (\text{Sin}[\pi - \Delta\tau + \Delta(t - \tau)], \text{Cos}[\pi - \Delta\tau + \Delta(t - \tau)], 0), \quad (3.30)$$

and therefore, after a total time 2τ , the system will be back in its initial state and an echo will be emitted. Another method of realigning the dipoles is using the atomic frequency comb (AFC) scheme. Here each dipole is separated in frequency by Δ and, therefore, the system after initial excitation² will be given by

$$\rho_n(t, \Delta) = (\text{Sin}[n\Delta t], \text{Cos}[n\Delta t], 0), \quad (3.31)$$

where ρ_n is the Bloch vector for the n th dipole. Now it can be seen that, after a time $t_s = 2\pi/\Delta$ all dipoles will again realign. Finally, there is the controlled reversible inhomogeneous broadening (CRIB)/gradient echo memory (GEM) method whereby, at time τ , we reverse the detunings of all dipoles so that the Bloch vector will be

$$\rho(t > \tau, \Delta) = (\text{Sin}[-\Delta(t - \tau) + \Delta t], \text{Cos}[-\Delta(t - \tau) + \Delta t], 0). \quad (3.32)$$

Again, it can be seen that after time 2τ the dipoles will realign.

3.4.2 Two-Level Gradient Echo Memory

We will discuss the two-level GEM scheme before extending to three-levels. As mentioned in the introduction, GEM relies on a linear frequency gradient η placed along the length of the memory L , giving a memory bandwidth

$$\mathcal{B}_{GEM} = \eta L. \quad (3.33)$$

The frequency-detuning for an atom at position z is determined by the gradient, i.e., $\Delta(t) = \eta(t)z$, and therefore a reversal of the gradient $\eta \rightarrow -\eta$ will lead to a reversal of Δ about the centre of the ensemble and the rephasing discussed in the previous section. To mathematically analyse the two-level GEM behaviour we use Eq.s 2.74 and 2.80 and relabel $\hat{\sigma}_{13} \rightarrow \hat{\sigma}_{12}$ to give

$$\frac{\partial}{\partial t} \hat{\sigma}_{12} = -(\gamma/2 + i\eta(t)z) \hat{\sigma}_{12} + ig\hat{\mathcal{E}}(\hat{\sigma}_{11} - \hat{\sigma}_{22}) \quad (3.34)$$

$$\left(\frac{d}{dt} + c\frac{\partial}{\partial z}\right) \hat{\mathcal{E}} = igN\hat{\sigma}_{12}. \quad (3.35)$$

In the weak probe regime we can assume $\langle \hat{\sigma}_{11} \rangle \approx 1$ and, moving into a reference frame moving at the speed of light while Fourier transforming the operators as per Eq. 3.11, we

²We continue to assume excitation by a $\pi/2$ -pulse here for ease. However, AFC and GEM both work in the weak probe regime where $\mathcal{A} \ll \pi/2$.




Figure 3.2: GEM Bandwidth, as a function of probe detuning Δ using Eq. 3.40 with $\text{OD}_{2l} = 0.3$ for an initial normalised bandwidth $\eta_o L$ for gradients (i) η_o , (ii) $2\eta_o$ and (iii) $3\eta_o$.

obtain the following equation for the propagation of $\hat{\mathcal{E}}$

$$\frac{\partial}{\partial z} \hat{\mathcal{E}}(z, \omega) = \frac{-g^2 N}{c} \hat{\mathcal{E}}(z, \omega) \frac{1}{\gamma - i(\eta z + \omega)}. \quad (3.36)$$

This can be solved to give

$$\hat{\mathcal{E}}(z, \omega) = \hat{\mathcal{E}}(-z_o, \omega) \mathcal{T}(z, \omega) \quad (3.37)$$

$$\mathcal{T}(z, \omega) = \left[\frac{\gamma - i(\eta z + \omega)}{\gamma - i(\omega - \eta z_o)} \right]^{i\text{OD}}, \quad (3.38)$$

where $z_o = -L/2$ and the two-level broadened optical depth is given by

$$\text{OD}_{2l} = \frac{g^2 N}{c\eta}. \quad (3.39)$$

Eq.s 3.37-3.38 contain a discontinuity for $\gamma = 0$. A correction for this, found in Ref. [186], gives $\hat{\mathcal{E}}$ at the end of the memory (now moving to a memory from $z = 0 \rightarrow L$ for ease of notation) as

$$\hat{\mathcal{E}}(L, \omega) = \hat{\mathcal{E}}(0, \omega) \text{Exp} \left[-\pi \text{OD}_{2l} \text{Hs} \left[\frac{\omega - \omega_o}{\eta L} \right] \right], \quad (3.40)$$

where $\text{Hs}[\omega]$ is the heavy-side step function (1 for $-0.5 \leq \omega \leq 0.5$ and 0 otherwise) and ω_o is the centre frequency. This is plotted in Fig. 3.2 for a range of bandwidths. As can be seen, as the gradient is increased the bandwidth increases (as per Eq. 3.33) while the absorption decreases (as per Eq. 3.40). If the pulse bandwidth fits within the memory bandwidth we can ignore $\text{Hs}[\omega]$ and see that the shape of the pulse will be preserved, with an amount $\text{Exp}[-\pi \text{OD}_{2l}]$ leaking through the memory. If the pulse bandwidth does not fit well within the memory bandwidth there will be large dispersion at the edges of the memory where the absorption changes sharply, as we saw with EIT earlier. We can define the efficiency of the storage process (remembering efficiency $\propto |\hat{\mathcal{E}}|^2$) to be

$$\varepsilon_s = 1 - \text{Exp}[-2\pi \text{OD}_{2l}]. \quad (3.41)$$

As GEM is a time-reversal process, the recall efficiency will be the same as the storage efficiency, giving a total efficiency of

$$\varepsilon_t = (1 - \text{Exp}[-2\pi\text{OD}_{2l}])^2. \quad (3.42)$$

For this reason, GEM can be conceptually represented as a beam-splitter process [186], with transmittivity $\text{Exp}[-\pi\text{OD}_{2l}]$ and can therefore be used for interference phenomena, as demonstrated experimentally by G. T. Campbell et al. in 2012 [10]. It can be seen that as $\text{OD}_{2l} > 1$ the total efficiency will approach unity, not limited by reabsorption issues in the forwards direction due to the monotonicity of the gradient.

There is also a phase shift associated with GEM storage, stemming from Eq. 3.38 with [224]

$$\Delta\phi(t) = \text{OD}_{2l} \ln \left[\text{OD}_{2l} + \frac{\mathcal{B}}{2}t \right], \quad (3.43)$$

with t being the storage time, assuming the pulse entered the memory at $t = 0$. This phase shift can be split into three terms, a constant phase shift (first term in the logarithm), a constant frequency shift and a frequency chirp. The latter two being due to the linear and nonlinear components of the second term in the logarithm. The frequency shift is proportional to $1/\mathcal{B}\tau$ and, as long as the storage time is much longer than the inverse of the memory bandwidth, and the OD is not too high, the frequency chirp will be negligible. This is the case in the experimental work presented in this thesis.

As GEM is a frequency-encoding memory we will not encounter the same pulse-bandwidth/length trade-off that was shown to exist for EIT. However, with the optical depth being inversely proportional to the bandwidth, increasing the bandwidth will lower the efficiency of the memory. It has been found that the multi-mode capacity of GEM scales with OD_{2l} [140], a factor of $\sqrt{\text{OD}_{2l}}$ better than EIT.

3.4.3 GEM Polariton

When discussing EIT we saw that defining a polariton, a mixture of optical and atomic modes, helped with determining the properties of the memory. We can create a similar polariton with GEM. To do so, however, involves taking a plane-wave decomposition of the optical and atomic fields by using a spatial Fourier transform of Eq. 3.34 and 3.35 to obtain [225]

$$\frac{\partial}{\partial t} \hat{\sigma}_{12} = -i\eta z \hat{\sigma}_{12} + ig\hat{\mathcal{E}} \quad (3.44)$$

$$\hat{\mathcal{E}}(k, t) = \frac{gN}{ck} \hat{\sigma}_{12}, \quad (3.45)$$

where we have ignored decoherence terms. From these relations we can obtain the following equation for the atomic polarisation

$$\left(\frac{\partial}{\partial t} + \eta \frac{\partial}{\partial k} - \frac{ig^2N}{ck} \right) \hat{\sigma}_{12} = 0. \quad (3.46)$$

This, plus a corresponding equation for $\hat{\mathcal{E}}$, can be solved by creating polariton states of the form

$$\hat{\Psi}_{\pm}(k, t) = k\hat{\mathcal{E}}(k, t) \pm \frac{gN}{c} \hat{\sigma}_{12}(k, t), \quad (3.47)$$

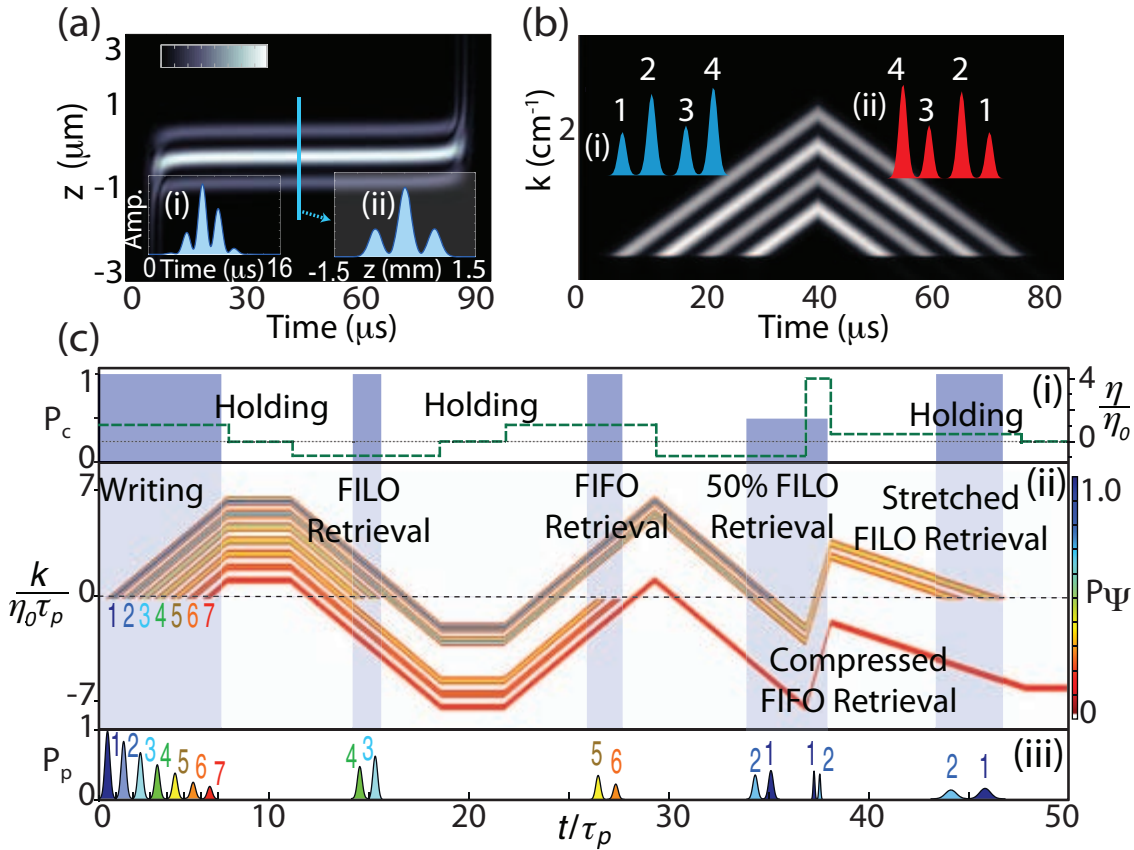


Figure 3.3: GEM Polaritons. (a) Plot of atomic polarisation (see legend) as a function of position inside memory ($-3 \mu\text{m} < z < 3 \mu\text{m}$) during storage. Insets show (i) temporal profile of pulse, (ii) cross-section of atomic excitation (which is the Fourier transform of the pulse). Adapted from Ref. [225]. (b) Absolute value of polariton for "normal" first-in-last-out storage. (i) Temporal profile on input pulses (blue) and (ii) output pulses (red). Adapted from Ref. [9]. (c) Pulse resequencing using Λ -GEM. (i) Shows normalised coupling field power (left hand scale shaded) and size of gradient (right hand scale, green dashed line), (ii) shows polariton evolution (see scale on right), and (iii) shows temporal profile of input pulses and echoes. Adapted from Ref. [9].

solving the equations

$$\hat{\Psi}_- = 0 \quad (3.48)$$

$$\left(\frac{\partial}{\partial t} + \eta \frac{\partial}{\partial k} - \frac{ig^2 N}{ck} \right) \hat{\Psi}_+ = 0, \quad (3.49)$$

where the solution for the negative state can be seen from Eq. 3.45.

The first piece of information we can take from the polaritonic description is the frequency-encoding nature of GEM storage, gleaned from the fact that the polariton is in k -space rather than z -space. This, in turn, means that the pulse is stored as its spatial Fourier transform along the memory. This is shown in Fig. 3.3(a). The polariton propagates in k -space³, starting at $k = 0$, at a rate determined by η . As the polariton travels to higher values of k , the light field component decreases due to dephasing of the dipoles discussed in Sec. 2.5.4. Again it can be seen that, if we reverse η , we will reverse the direction of propagation of the polariton in k -space, leading to re-emission at $k = 0$ (corresponding to realignment of the dipoles).

If we recall with a steeper (shallower) gradient the polariton will pass through $k = 0$ for a shorter (longer) time, leading to pulse compression (stretching). There will be a corresponding change in optical depth and efficiency via Eq.s 3.39 and 3.42. If we were to turn off the gradient after the pulse had entered the memory, then the polariton would stop evolving in k -space. This is illustrated by the ‘‘Hold’’ labels in Fig. 3.3(c). It can also be seen from the evolution of the polariton in Fig. 3.3(b) that GEM acts as a first-in-last-out memory, reflecting its time-reversal nature. With non-unit retrieval efficiencies, left-over energy remains in the memory and this (much smaller) polariton can be reversed again to recall in a first-in-first-out manner. However, if we were to have a ‘‘knob’’ on the memory that would allow us to suppress emission at $k = 0$ on the first pass, then we would be able to achieve high efficiency first-in-first-out recall, as well as the ability to resequence pulses.

3.4.4 Stepping Up a Level - Three-Level GEM

There are a few reasons for moving from a two-level version of GEM to a three-level, or Λ , version. The first of these is that the excited-state decay rate γ depends on the coupling strength between the two levels, as does the optical depths and therefore efficiency. This means that we can either have long storage times, or high efficiencies, but not both. Secondly, with a three-level system we have an extra ‘‘knob’’ to play with (see below).

Let us start with a Λ -system similar to that for EIT, except with a large excited-state one-photon detuning Δ , shown in Fig. 3.4(a). We will again use Eq.s 2.70-2.75 and Eq. 2.80 while making the adiabatic approximation ($\partial\hat{\sigma}_{13}/\partial t \ll \gamma$), which is equivalent to $1/T \ll \gamma$, where T is the fastest time-scale of the system, and $\Delta \gg \gamma$. For an ensemble memory, these conditions become $\text{OD}_{2l} \gg 1/T\gamma$ and $\text{OD}_{2l} \ll \Delta/\gamma$, leading to $\Omega_c/\Delta \ll 1$ [214]. With these assumptions, while still assuming a weak probe and moving to a frame travelling at the speed of light, we find

$$\frac{\partial}{\partial t} \hat{\sigma}_{12} = - \left(\gamma_o + i\eta(t)z - i\frac{\Omega_c^2}{\Delta} \right) \hat{\sigma}_{12} - i\frac{g\Omega_c \hat{\mathcal{E}}}{\Delta} \quad (3.50)$$

³Technically a new polariton, with its own commutation relations, is created at each new time interval δt [225].




Figure 3.4: Two-Level, Three-Level Equivalence. (a) Three-level atom used for Λ -GEM, with: probe field $\hat{\mathcal{E}}$; atom-light coupling strength g ; one-photon detuning Δ ; two-photon detuning δ depending on the position of the atom in the ensemble (z) and the size of the gradient η ; coupling field with Rabi frequency Ω_c ; and decoherence rate γ_o . (b) Equivalent two-level atom, with: excited-state decay rate γ_o ; detuning δ ; and modified coupling strength $g' = (\Omega_c/\Delta)g$.

$$\frac{\partial}{\partial z}\hat{\mathcal{E}} = \frac{igN}{c}\frac{\Omega_c}{\Delta}\hat{\sigma}_{12} + \frac{ig^2N}{c}\hat{\mathcal{E}}, \quad (3.51)$$

where the two-photon detuning $\delta(z, t) = \eta(t)z$. The Ω_c^2/Δ term in Eq. 3.50 is due to the ac Stark shift from the coupling field. If we were to choose a coupling field frequency to compensate for this term, and move into a frame where the speed of light in the medium is normalised by $ig^2N\hat{\mathcal{E}}/c$ then these equations become

$$\frac{\partial}{\partial t}\hat{\sigma}_{12} = -(\gamma_o + i\delta(z, t))\hat{\sigma}_{12} - i\frac{g\Omega_c\hat{\mathcal{E}}}{\Delta} \quad (3.52)$$

$$\frac{\partial}{\partial z}\hat{\mathcal{E}} = \frac{igN}{c}\frac{\Omega_c}{\Delta}\hat{\sigma}_{12}, \quad (3.53)$$

formally equivalent to Eqs. 3.34 and 3.35 if $g \rightarrow g\Omega_c/\Delta$ and $\gamma \rightarrow \gamma_o$. This is shown in Fig. 3.4(b). The new OD for the system will then be

$$\text{OD}_{3l} = \frac{g^2N}{c\eta} \left(\frac{\Omega_c}{\Delta}\right)^2. \quad (3.54)$$

With Λ -GEM we can therefore alter the OD during the storage process by changing the coupling field power. If we were to turn off the coupling field completely we would have an OD of zero and therefore no emission would occur, even if the polariton passed through $k = 0$. This is the extra “knob” we use to allow for pulse sequencing and beam-splitting operations in Λ -GEM (see Fig. 3.3(c)) [9]. By altering the coupling field frequency we can also change the frequency of the output pulse, allowing us to compensate for the constant frequency shift discussed in Sec. 3.4.2.

Another benefit of moving from two to three levels is that the decay is now limited by $\gamma_o \ll \gamma$. Experimentally this opens up a range of readily available storage medium options, with large numbers of addressable atoms at convenient laser wavelengths, that

would otherwise be unusable due to their fast excited-state decay rates (for instance, see App. A). One disadvantage of using three levels is that, as $\Omega_c/\Delta \ll 1$, $\text{OD}_{3l} \ll \text{OD}_{2l}$ assuming the same coupling strength.

3.5 Conclusions

In this chapter we have presented the theory behind the two memory techniques we will use in this thesis: electromagnetically-induced transparency; and the three-level gradient echo memory. We showed how an EIT-based memory stores light by reducing its group velocity to zero and back again. GEM, in comparison, is based around reversing the initial absorption process. We discussed some of the properties of Λ -GEM that make it a promising quantum memory candidate: high efficiencies for large ODs; the coherence time being determined by the ground-state decoherence rate; the frequency-encoding nature of the storage process that allows for multi-mode storage; as well as the ability to manipulate the stored information using the gradient and coupling field.

Aside
Digital Locking

A Digital Locking System for use on Quantum Optics Experiments

HAL: "I am putting myself to the fullest possible use, which is all I think that any conscious entity can ever hope to do."

- 2001: *A Space Odyssey*, MGM, April 1968, Movie.

4.1 Introduction

Examples of the use of digital control algorithms, rather than analog electronics, to lock the frequency of single [226] and multiple [227] lasers date back to 1998. This progression has been motivated by the fact that a digital system can offer distinct advantages over its analog counterpart. These include cost savings (as features can be programmed rather than purchased), as well as space savings (as most of the features can be included in the code they only take up as much space as the hardware needed to run it). Mostly, however, the advantage of digital locking over analog systems lies in their flexibility: the ability to change the functionality of the system by altering the control code rather than having to purchase new equipment or physically modify old equipment. This makes a digital system easily reproducible, as well as having the ability to add complex logic that would be difficult to include with analog circuits alone (for examples see Ref. [228]). Using the added power of digital control has led many groups to develop novel methods for frequency stabilising lasers (see, for example, Ref.s [229, 230, 231, 232, 233, 234]).

Complex logic and functionality are not entirely the domain of digital technology, with both G. J. Koch [235] and S. K. Lee et al. [236] having developed methods using analog electronics to automatically lock lasers to gas absorption lines and a Fabry-Perot cavity respectively, with S K. Lee's group being able to unlock and relock the cavity at will. However, with digital control this technology can be improved with, for instance, L. Dong et al. [237] using LabVIEW and an alternate method of scanning and noting of frequencies to allow their system to lock to any selected gas absorption line.

Quantum optics is one branch of science that can benefit from digitisation and computer control. One component common to many quantum optics experiments is the cavity. For instance, mode-cleaner-style ring cavities to select only one spatial mode of an input field, or Fabry-Perot cavities to select a specific frequency. Cavities are not just used for filtering, however, and can also be an integral part of an experiment. For instance, cavities containing non-linear optical crystals (known as optical parametric oscillators - OPOs) are used to produce entanglement [26] and squeezing [238, 239], two powerful quantum resources. These effects are extremely sensitive to noise or loss in the system. Therefore

the properties of the system used to lock an OPO must be stringently monitored. As well as the use of cavities, another issue of concern is the relative phase between various fields. For instance, homodyne detection using a local oscillator [240, 241].

Quantum optics experiments can require many locking loops and sophisticated data acquisition, and therefore these experimental systems can become highly complex. As an example, in 2003 W. Bowen et al. demonstrated the teleportation of the quadrature amplitudes of two light fields [242]. To achieve this the experiment required a frequency doubling cavity, a high finesse ring cavity used to seed a pair of OPOs, as well as two homodyne detector set-ups requiring phase control.

More recently, another example of the complexity of quantum optics experiments is the work carried out by M. Yukawa et al. in 2008 to generate four-mode cluster states for use in quantum computing [243]. This experiment required a frequency doubler, which was used to pump four OPOs, and four homodyne set-ups were required for the measurements giving a total of nine independent locking loops, each requiring long-term stability. From these examples it can be seen that a digital control system for quantum optics experiments would need to be extendible to many locks, flexible in the style of locks it can control, as well as be as effective, or better, than current analog controllers in reducing noise.

As more complex experiments are developed to investigate further into the quantum realm, the practicality of digital control will become more apparent. This is because digital control allows for integration of all relevant information about the system to be accessed, and controlled, from one location. This, in turn, allows the system to take into account the sequential nature of the locks comprising the experiment. Also, the complete system information can be used for conditional data acquisition, to ensure that data is only recorded when the system is in the desired state. As well as this, digital interfaces for triggering experimental components are a fast, user-friendly method to control timings for experiments (these will be discussed more in the next chapter).

In this chapter we present a code designed for quantum optics experiments using field programmable gate arrays (FPGAs) programmed with National Instruments (NI) LabVIEW software. This code is freely available to be downloaded and to be modified, as both a pedagogical tool, and to allow individual users to tailor it to their specific needs (see App. B).

The remainder of this chapter is structured as follows: an introduction to control theory is presented in Sec. 4.2, followed by an overview of the hardware (Sec. 4.3.1) and software features (Sec. 4.3.2) of the code. In Sec. 4.4 we show how the inbuilt locking analysis tools, including a white-noise network analyser, can be used to help optimise individual locks, and verify the long term stability of the digital system against an analog proportional-integral controller. Finally, in Sec. 4.5, we present an example of the benefits of digital locking for quantum optics by applying the code to a specific experiment used to characterise optical Schrödinger cat states.

Most of the work presented in this chapter has been published in the journal articles:

- B. M. Sparkes, H. M. Chrzanowski, D. P. Parrain, B. C. Buchler, P. K. Lam, and T. Symul, *A Scalable, Self-Analyzing Digital Locking System for use on Quantum Optics Experiments*, Review of Scientific Instruments **82**, 0751133 (2011)
- H. M. Chrzanowski, J. Bernu, B. M. Sparkes, B. Hage, A. P. Lund, T. C. Ralph, P. K. Lam, and T. Symul, *Photon-number discrimination without a photon counter*

and its applications to reconstructing non-Gaussian states, Physical Review A **84**, 050302(R) (2011).

4.2 Control Theory

This section provides an overview of the relevant control theory, for a more in depth description see, for example, Ref. [244]. Figure 4.1(a) shows the basics of a closed-loop feedback system. A plant - the system to be controlled - having a transfer function (i.e., frequency response) $G(\omega)$ produces some form of error signal $Y(\omega)$. This is passed to the controller, having a transfer function $H(\omega)$, used to suppress the noise added to the system $X(\omega)$.

The system, or closed-loop, transfer function $T(\omega)$ can be determined from $G(\omega)$ and $H(\omega)$ by using the definition of a transfer function, namely

$$T(\omega) = \frac{Y(\omega)}{X(\omega)}. \quad (4.1)$$

From Fig. 4.1(a) and the fact that, for transfer functions, components in series multiply, we can see that the actuating signal $E(\omega)$ depends on the controlling signal and the input noise via

$$E(\omega) = X(\omega) - R(\omega) = X(\omega) - H(\omega) \cdot Y(\omega), \quad (4.2)$$

and also that the output signal depends on the actuating signal and the plant transfer function via

$$Y(\omega) = E(\omega) \cdot G(\omega). \quad (4.3)$$

Substituting Eq. 4.2 into Eq. 4.3 gives

$$\begin{aligned} Y(\omega) &= [X(\omega) - H(\omega) \cdot Y(\omega)] \cdot G(\omega) \\ [1 + G(\omega) \cdot H(\omega)] Y(\omega) &= X(\omega) \cdot G(\omega) \\ \frac{Y(\omega)}{X(\omega)} &= \frac{G(\omega)}{1 + G(\omega) \cdot H(\omega)}, \end{aligned} \quad (4.4)$$

which is the definition of $T(\omega)$ given by Eq. 4.1.

For quantum optics experiments, as with all locking servo applications, we are interested in three properties of the controller: how well it can suppress noise at low frequencies; what frequency range it can achieve this over; and how stable the lock is. These properties can be gleaned from the equation above. Firstly, the larger $H(\omega)$ the smaller $T(\omega)$ will be. This means that the noise entering the system $E(\omega)$ is being suppressed, with more suppression occurring as $H(\omega)$ increases, and the amount of suppression determined by the difference between $T(\omega)$ and $G(\omega)$. Conversely, as $H(\omega) \rightarrow 0$, $T(\omega) \rightarrow G(\omega)$ and therefore there will be no noise suppression. The frequencies for which noise suppression occurs define the bandwidth of the system - the range of frequency for which the system can compensate for the input noise. This is normally defined mathematically as the point where $|GH(\omega_B)|$, known as the open-loop transfer function, equals one. This is also known as the unity gain frequency ω_B . Finally, as $G(\omega) \cdot H(\omega)$ approaches -1 , $T(\omega) \rightarrow \infty$ and therefore the system becomes unstable. The stability of the system is usually measured

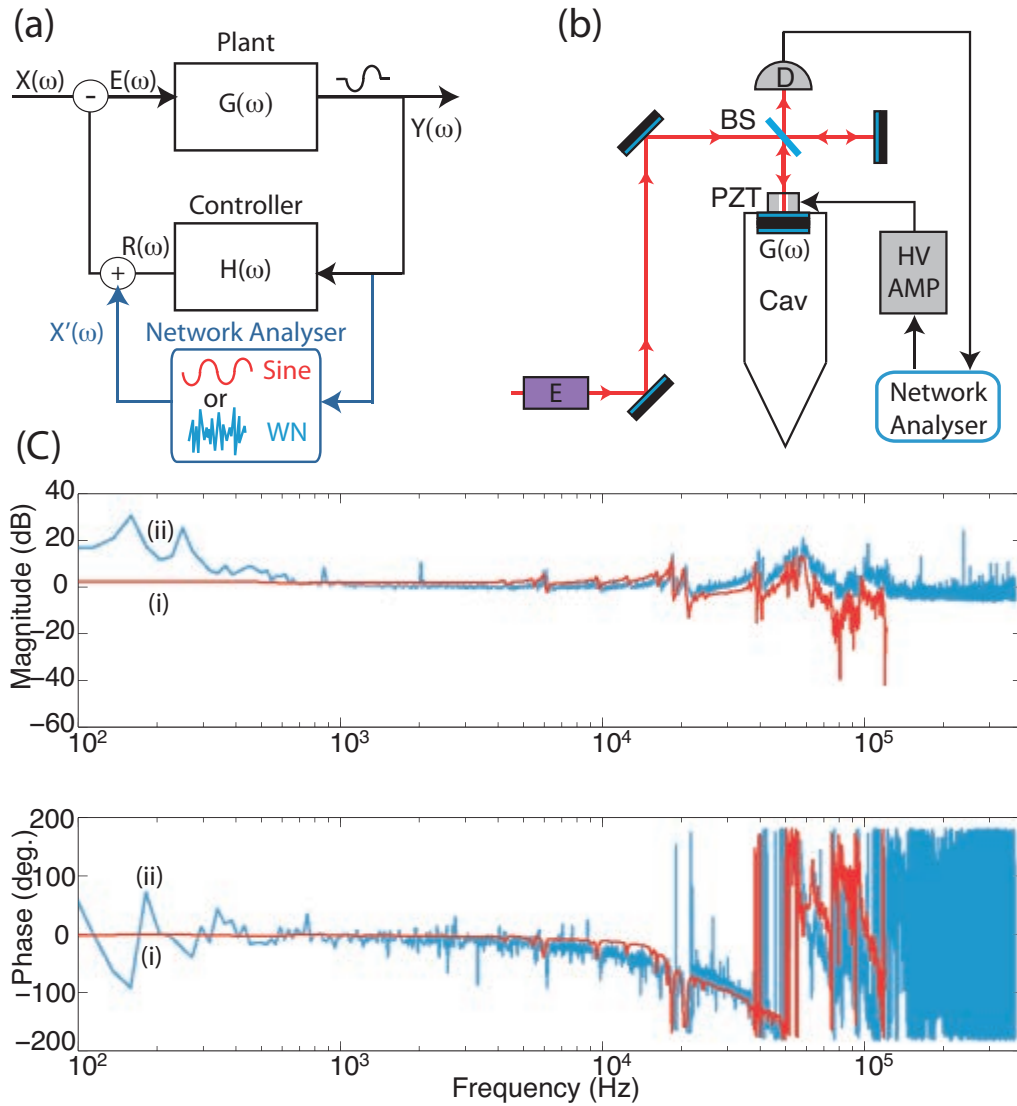


Figure 4.1: Feedback Control System. (a) Block diagram of a basic feedback control system. Here $G(\omega)$ and $H(\omega)$ are the transfer functions of the plant and controller respectively, $X(\omega)$ is the input noise, $Y(\omega)$ is the output of the system, $R(\omega)$ is the feedback signal, and $E(\omega)$ is the actuating error signal. Also shown is the placement of a network analyser, with $X'(\omega)$ the extra noise added to the system in the form of a swept sine wave (Sine) or white noise (WN). (b) Experimental set-up for measuring $G(\omega)$ for a ring cavity with: Cav - ring cavity; E - electro-optic modulator; BS - 50:50 beam-splitter; PZT - piezo-electric transducer; D - detector; and HV AMP - high voltage amplifier. (c) Bode plot of the transfer function of a ring cavity for (i) a network analyser (model MS4630B from Anritsu) and (ii) a white noise generator included in the code.

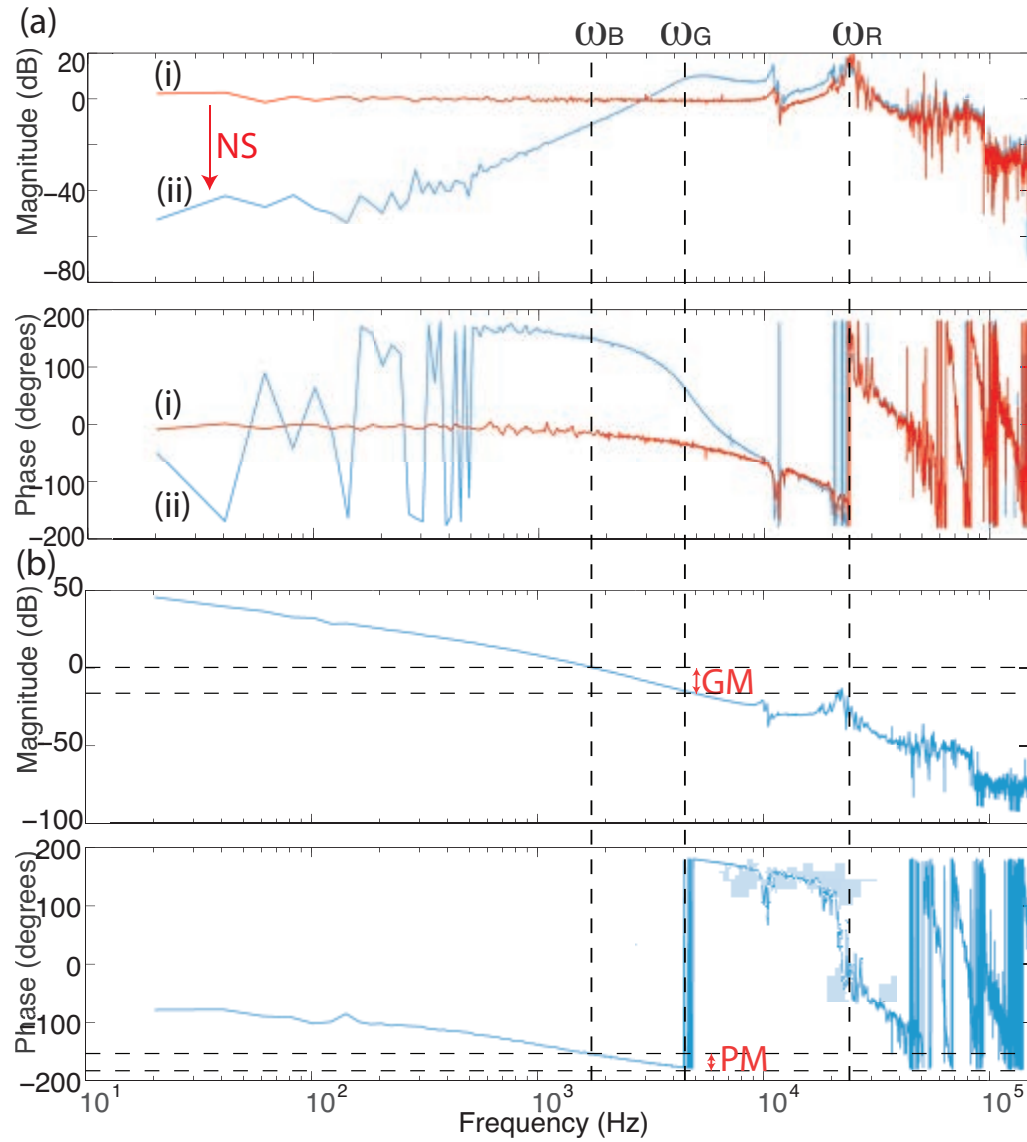


Figure 4.2: Controller Properties. (a) The Bode plot of (i) $G(\omega)$ and (ii) $T(\omega)$ for the system shown in Fig. 4.1 and (b) the corresponding open-loop transfer function with: ω_B - system bandwidth; ω_G - frequency satisfying $\angle GH(\omega_G) = \pi$; ω_R - system resonance frequency; NS - low-frequency noise suppression; PM - phase margin; GM - gain margin.

using the Nyquist stability criterion [245] via the phase margin (PM) defined by

$$PM = \pi + \angle GH(\omega_B). \quad (4.5)$$

That is, the clearance of the phase of $GH(\omega_B)$ from $-\pi$. Conversely, if we define ω_G to be the frequency at which $\angle GH(\omega_G) = \pi$, then we can define another stability criterion, the gain margin (GM), as

$$GM = \frac{1}{|GH(\omega_G)|}. \quad (4.6)$$

These properties are illustrated in Fig. 4.2.

Ideally we would like to make the bandwidth, as well as the size of $H(\omega)$, as large as possible for maximum noise suppression over all frequencies. However, most systems have some resonance, at frequency ω_R , above which it is not possible to suppress noise. Delay within in a feedback system will also limit the frequency response of a controller. The frequency response of each system to be locked should, therefore, be investigated to allow for optimal locking. One way of investigating the system transfer function is to use a network analyser to add known noise $X'(\omega) \gg X(\omega)$ to the system and measuring the frequency response. For commercial network analysers (such as the Anritsu model MS4630B used here) this extra noise is usually in the form of a swept sine wave, though white noise - that is noise with the same amplitude at all frequencies - can also be used. The placement of a network analyser to measure the system transfer function is shown in Fig. 4.1(a). We investigated the unlocked transfer function for a ring cavity, used for mode-cleaning applications, where we drove a high-voltage (HV) amplifier that, in turn, drove a piezo-electric transducer (PZT) on the back mirror of the cavity (see figure) to change the cavity's length. To take data for an unlocked system it was necessary to build a Michelson interferometer at the rear of our cavity, using the back cavity mirror as one of the Michelson arms. This is shown in Fig. 4.1(b).

A Bode plot shows both the magnitude ($10 \cdot \log_{10} |T(\omega)|$) and phase delay ($\angle [T(\omega)]$). Examples of Bode plots for the ring cavity system are shown in Fig. 4.1(c). Note the flatness of $G(\omega)$ at frequencies below 10 kHz and compare this to the suppression of noise shown in Fig. 4.2(a). It can also be seen that the phase delay $\Delta\phi$ (which is determined by the time delay Δt via $\Delta\phi = \omega \cdot \Delta t$) reaches $-\pi$ at approximately $\omega_R = 50$ kHz. Once $\Delta\phi > -\pi$ feedback control will no longer work as the controlling and output signals will be π rad (or 180°) out of phase and therefore when the feedback signal should be positive it will be negative and vica-versa, causing the system to resonate. This limits the controller bandwidth to some frequency below the plant resonance.

In practice we must take into account the fact that $G(\omega)$ is composed of the responses of the cavity, the photo-diode detector, the analog-to-digital converter (ADC), the digital demodulation and filtering, as well as the digital-to-analog converter (DAC), the HV amplifier and the PZT, not to mention the time delay due to the cables and through the digital controller. However, apart from the PZT-cavity system $G'(\omega)$, the responses of these components in the frequency range we are interested in (< 100 kHz) should be approximately flat and therefore only add constant gain to the system (i.e. $G(\omega) \approx \mathcal{K} \cdot G'(\omega)$ where \mathcal{K} is a constant).

Another measure of the efficacy of a lock is the deviation of the error signal from its desired value over time. This is usually measured using the root mean squared (RMS)

method, which is expressed mathematically as

$$\text{RMS} = \sqrt{\frac{\sum_{i=1}^n (S_i - S_{des})^2}{n}}. \quad (4.7)$$

Here S_i is the value of the error signal at a particular time i , S_{des} is the desired value of the error signal (usually zero), and n is the number of points the RMS was measured over.

4.3 The Digital Control System

The digital control system presented here is based on the Pound-Drever-Hall (PDH) locking technique [246, 247], presented in App. C. This technique uses phase modulation of the laser frequency to produce an anti-symmetric error signal (see Fig. 4.3(a) inset) that is then fed back to the actuating mechanism to keep the system locked to a desired position. The code allows for automatic, sequential locking, as well as analysis of the locks comprising the experiment to facilitate optimizing procedures and these will all be discussed here. More details on the code can be found in App. B.

4.3.1 The Hardware

Figure 4.3(a) shows the basic setup for a single lock (in this case a mode-cleaner cavity). A frequency generation (FG) signal at 80 MHz (NI PXI-5404) is split in two (MiniCircuits splitter ZSC-2-1), with one output being sent to a clock generator board (CGB - AD9959 from Analog Devices) and one to clock a high-speed analog-to-digital convertor (HS ADC - AD9460BSVZ-80 from Analog Devices at 80 M Samples s^{-1}). The output from the CGB is controlled via an FPGA card (NI PXI-7852R), and sends a sine wave of a desired frequency to some form of modulation device - here an electro-optic modulator (EOM - Newport free-space model 4002-M) - to modulate the phase of the input field.

Once the modulated signal passes through the cavity, part of it is reflected on to a photo-detector that has two output ports - one ac-coupled and one dc-coupled. The ac component, containing the modulation signal, is sent to the HS ADC. The dc component, if there is one, is sent to a low-speed analog input (AIP \approx 1 M Samples s^{-1}) of the FPGA. The controller algorithm is discussed in Sec. 4.3.2, with the output signal sent through a low-speed analog output (AOP), via a high-voltage (HV) amplifier to the piezo-electric transducer (PZT) controlling the round-trip cavity length.

This system has been designed with scalability in mind to cater to experiments with large numbers of locks. Using the same principle as above, up to eight locks can be implemented using two FPGAs, two CGs, one FG, and eight HS ADCs. This is achieved by splitting the 80 MHz FG signal as many times as necessary to send it to clock all HS ADCs and one signal to each CGB (both controlled by one of the FPGAs). Also, the 16-bit signals from pairs of HS ADCs are combined into one 32-bit signal before being sent to an FPGA. This is because the FPGA cards used here have only two 40-bit digital input/output (DIO) connections, and in this way four locks can be implemented on each FPGA. Each FPGA also has eight low-speed AIPs and eight low-speed AOPs, and another 16 DIO lines all located on a mixed IO connector. These are more than enough for the four dc input and four controller output signals required per FPGA, as well as extra digital lines to control the CGBs.

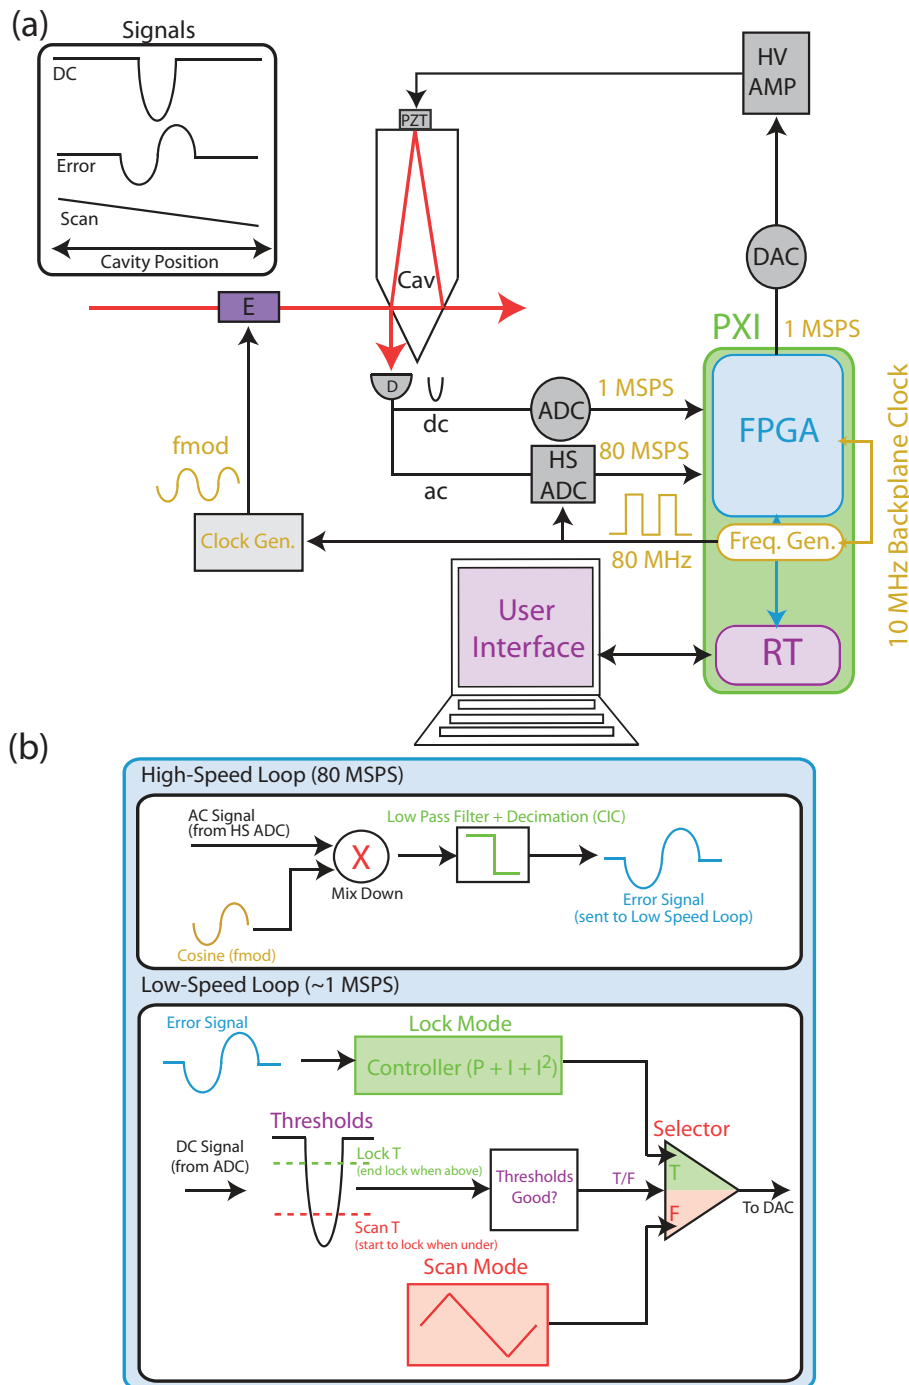


Figure 4.3: Digital Locking System Set-Up. (a) Digital locking hardware set-up for a mode-cleaner ring cavity (Cav) with: MSPS - MSamples s^{-1} ; PXI - PXI chassis; FPGA - field programmable gate array card; Freq. Gen. - frequency generator; RT - real time controller; AIP - analog input to FPGA (sampling frequency < 1 MSPS); AOP - analog output from FPGA (sampling frequency < 1 MSPS); HS ADC - high-speed analog-to-digital converter (80 MSPS); D - detector; Clock Gen. - digital clock generation board; f_{mod} - modulation frequency for lock; E - electro-optic modulator; PZT - piezo-electric transducer; HV AMP - high-voltage amplifier. Inset shows the different signals that are used for locking (DC - dc reflection signal; Error - Error signal; Scan - scan function) as a function of PZT position. (b) Depiction of LabVIEW code used to program FPGA cards. CIC - cascaded-integrator-comb filter; PII controller - proportional, integral, double integral controller; Scan T - scan threshold; Lock T - lock threshold; T - true; F- false.

(a)



(b)



Figure 4.4: Digital Locking Racks. (a) ADC rack with: *Clock* - 80 MHz input; inputs - $IP_{x.1}$ and $IP_{x.2}$; and combined outputs OP_x that is then sent to a DIO terminal on the PXI chassis. (b) Clock rack with: $CLK-IN$ - 80 MHz input; $CTRL$ - FPGA control input; and $DAC-x$ - frequency outputs.

Using this method up to 16 locks can be implemented on one PXI chassis (here a NI PXI-1042Q). We took the set-up one step further by having rack-mountable boxes designed for both the ADCs and CGBs. These are shown in Fig. 4.4

4.3.2 The Software

The code was developed using LabVIEW 2010 (32 bit). An overview of the FPGA code is shown in Fig. 4.3(b) and consists of two loops: a high-speed loop running at $80 \text{ M Samples s}^{-1}$; and a low-speed loop running at approximately $750 \text{ k Samples s}^{-1}$ (limited by the low-speed AIP/AOP). In the high speed loop the modulated signal from the HS ADC is demodulated with a cosine function generated using a look-up table at the same frequency as the modulation (i.e. f_{mod}). This passes through a cascaded-integrator-comb (CIC) filter [248] to produce an error signal that is then sent to the low-speed loop.

There are two main components to the low-speed loop: a lock mode (consisting of a P, I and I^2 - PII - controller, discussed in more detail in Sec. 4.3.3); and a scan mode (consisting of a saw tooth scan function). Once directed by the user to lock, the code will move from scan mode to lock mode once the dc signal drops below the user defined scan threshold (determined as a fraction of the peak size measured by the code, see Fig. 4.3(b)). The code will then remain in lock mode unless either the user disengages locking or the dc input rises above the lock threshold (set higher than the scan threshold). In both cases the system will revert to scan mode. In the latter case, however, the system will move back to lock mode once the dc input again drops below the lock threshold, unless the *Do Not Relock* option is engaged¹. If the lock has no dc component (i.e. a phase-style lock) the absolute value of the error signal can be used to set the thresholds instead. Alternatively, if desired, both thresholds can be set to zero and the user can manually lock the system as they would with an analog controller.

The low-speed loop includes a white noise generator, and this is used to measure the system transfer function as described in Sec. 4.2.

¹Useful for situations where two peaks are of approximately equal height, as we will discover in the next chapter

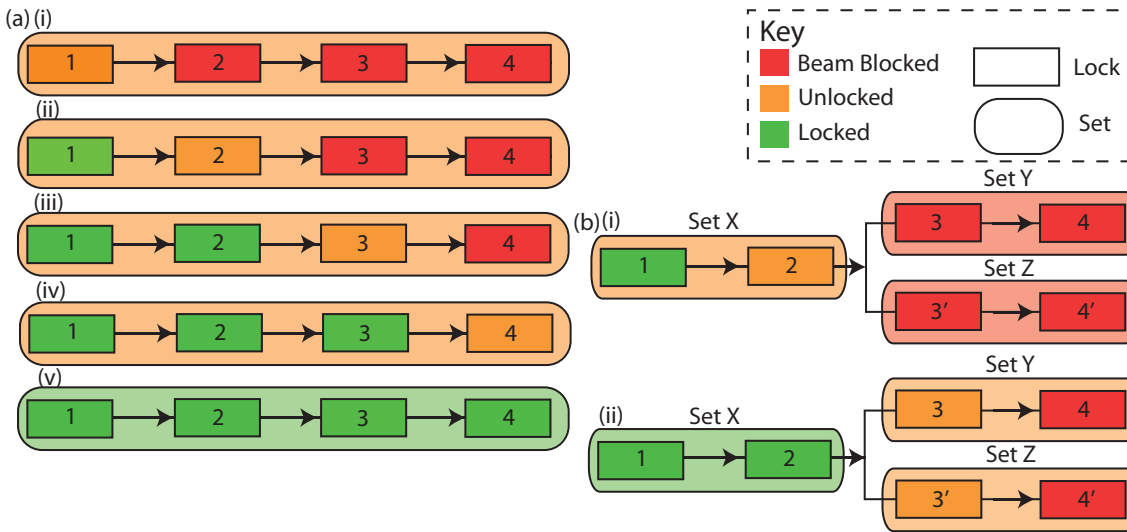


Figure 4.5: Sequential Locking. A schematic diagram of the principles behind the sequential locking logic developed for the code for (a) a single set of four sequential locks (arrows indicated dependencies), and (b) for two sets of two locks with sets Y and Z depending on Set X.

A master code runs on the odd-numbered FPGAs and a slave code runs on the even-numbered FPGAs. The codes differ in that the master code contains an internal PXI trigger for the FG (only on the first FPGA) and also runs the CGB programming, ensuring that the f_{mod} used for demodulation in the high-speed loop is the same frequency as that sent to drive the modulator, and with a fixed phase relation.

A LabVIEW Real Time (RT) code acts as the interface between the FPGAs and the user, as well as to allow communication between FPGAs, with all FPGA cards being controlled from the one RT code. The RT code contains logic that does not need to be included in the FPGA code, for instance a *Beam Blocked* function and a *lock Time-Out* function. If the beam is blocked before the lock while the code is running, then the minimum value for the dc signal will be set to zero. When it becomes unblocked the lock will more than likely no longer be on resonance and therefore the system will start scanning as discussed above. However, as the minimum registered value is now zero for both the scan and lock thresholds, the dc value at which the code will move into lock mode will be much lower than before meaning that, in most cases, the system cannot relock. To overcome this problem, if the user knows that they are going to block the light before the cavity, the *Beam Blocked* button can be activated. This freezes the cavity in its current position and doesn't register any new dc values until the button is released. Following from this, the idea for the *Time-Out* function is that, if the laser is blocked inadvertently or the power drifts over time (or any other event that would alter the maximum and minimum values), the program will try to re-lock for a certain amount of time (specified by the user). After this time, if it is unsuccessful, the minimum and maximum values are reset, allowing the system to lock again with new thresholds.

The ability for the RT code to communicate between FPGAs allows for sequential locking logic to be implemented across all FPGAs. Here the user can specify the order of the locks in their experiment, and which locks are dependent on which others. The locks are grouped into sets, where each lock in a set depends on the lock before it in a linear

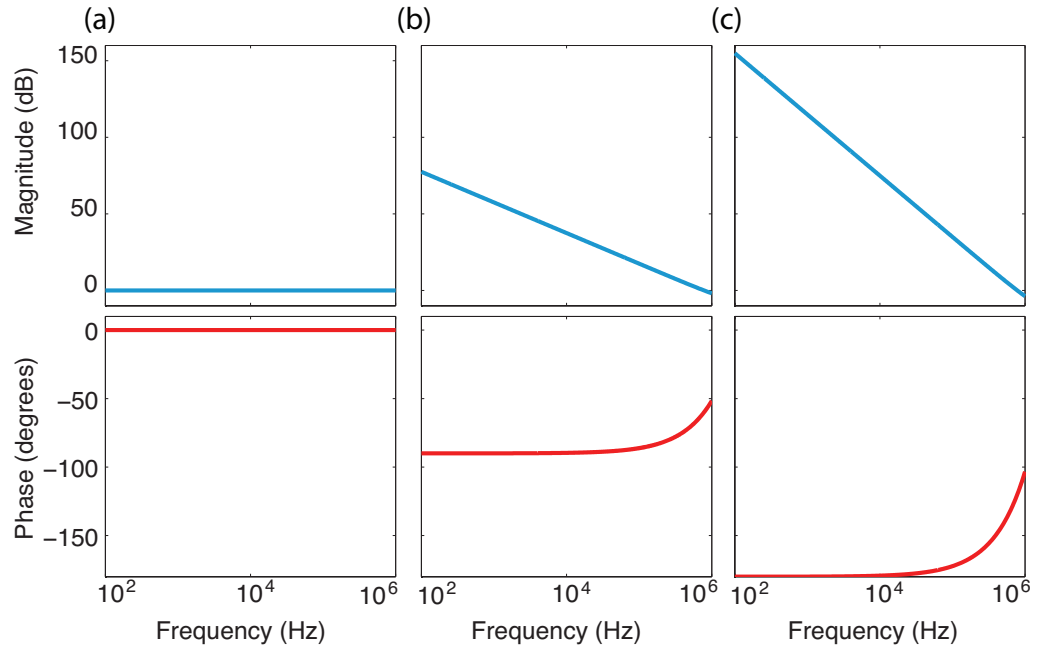


Figure 4.6: PII Controller. Bode plots for (a) proportional, (b) integral and (c) double integral controllers with $k_P = k_I = k_{II} = 1$.

fashion, and sets can depend on any other set, with multiple sets allowed to depend on the same set. This is illustrated in Fig. 4.5. If a lock drops then the controller signal for any locks that depend on the dropped lock will become *Beam Blocked* until the first lock is restored, when they proceed to relock one after the other.

The RT code also includes an inbuilt scope which can display information about any of the locks such as error signal, dc signal and scan signal in real-time on the user's computer monitor.

4.3.3 The Controller

The simplest controller design imaginable is to multiply the error signal by some constant gain k_P (either positive or negative), known as a proportional (P) controller. This does not allow much room for optimisation, however, with the same controller gain at all frequencies. For instance, this method limits the amount of noise suppression at low frequencies as, after a critical value of k_P , any additional proportional gain will lead the system to resonate. This is why usually a proportional-integral-derivative (PID) controller is used, with the integral (I) component (with gain k_I) for extra gain at low frequencies and the derivative (D) component doing the same at high frequencies.

In this experiment we are more concerned with noise at low frequencies rather than those above the cavity resonance (with ω_R on the order of 10 kHz) and therefore we decided to use a proportional-integral-double integral (PII) controller² with the double-integrator

²Though of course, as stated previously in this chapter, one of the advantages of digital locking is the flexibility to choose any desired controller algorithm.

(II) having gain k_{II} . For a continuous system the PII controller has the following transfer function [244]

$$H(s) = k_P + \frac{k_I}{s} + \frac{k_{II}}{s^2}, \quad (4.8)$$

where s is the Laplace variable that can be expanded in terms of the imaginary number i and frequency $f = \omega/2\pi$ to give $s = 1/if$. However, a digital system is discretised rather than continuous and, in this case, the transfer function becomes

$$H(z) = k_P + k_I \cdot \frac{z}{1+z^2} + k_{II} \cdot \left(\frac{z}{1+z^2} \right)^2, \quad (4.9)$$

where $z = \text{Exp}[if/f_s]$, with f_s being the sampling frequency of the controller system. Bode plots for the three different components of the PII controller are shown in Fig. 4.6. As already stated, the P component has the same gain at all frequencies, while the I and II have more gain at lower frequencies. It is also worth noting that both the I and II components have a phase delay associated with them ($\pi/2$ and π radians respectively) that can affect the stability of the system if their bandwidth extents close to the resonant frequency of the system. At higher frequencies, where the I and II components no longer suppress noise, this phase delay decreases.

For ease of use in situations where the ratio of the gains k_P , k_I and k_{II} is optimised but the system gain needs to be increased/decreased (for instance, due to an increase in error signal size), an extra controller gain parameter k_H has been added to the code to scale all other gains (i.e. $H_{act}(z) = k_H \cdot H(z)$).

4.4 Locking Analysis and Optimisation

In this section we present an experimental analysis of our locking system acting on both a ring cavity, and bow-tie OPO cavity containing a periodically poled potassium titanyl phosphate (PPKTP) crystal and co-resonant with light at 1064 and 532 nm.

4.4.1 Measuring System Responses

It is important, for both the optimisation and stability of a lock, to know the frequency response of the system in question. This normally requires a network analyser. However, these devices are expensive pieces of equipment (usually on the order of \$10,000 AUD) and have a minimum sine amplitude which can, in some cases, be enough to unlock the system being measured (even with extra external attenuation), adversely affecting the results. Another option, and one that is simple to implement in LabVIEW thanks to an inbuilt waveform generator, is to use white noise - a signal with the same amplitude at all frequencies. Not only is it easy and costs nothing to implement on top of LabVIEW and PXI set-up costs, but the output signal amplitude can be as low as $3 \cdot 10^{-4}$ V and, as an added extra, the code can be set to only take data when the system of interest is locked, and has been locked for a set amount of time to allow it to settle, removing the possibility of this affecting the results. One area where a network analyser has an advantage over white noise generation is resolution at low frequencies as, due to memory limitations on the FPGA cards, only approximately half a seconds worth of data can be taken at any one time. This can, however, be improved by averaging multiple traces, with 2^{10} traces of 2^{15}




Figure 4.7: Approximating G and Simulating H . (a) Approximation of $G(\omega)$ by using very small controller values and measured using a white noise generator. (b)(i) A comparison of $|H|$ calculated from Eq. 4.10, with G and T measured using LabVIEW (red), and (ii) $|H|$ simulated from Eq. 4.9 (blue).

points and $f_s \approx 747$ kHz used for all analysis below. A comparison of the two methods is shown in Fig. 4.1(b).

Both methods require timescales on the order of hundreds of seconds to obtain sufficient resolution at low frequencies. This proves to be a problem for measuring G for all but the most stable of cavities as, if the system is not very close to resonance, there will be no output signal (see inset Fig. 4.3(a)). One way around this, as mentioned previously, is to use a Michelson interferometer (Fig. 4.2(b)) as a signal incident on the detector will be a sine wave (and therefore the system will never be in a range where the output signal will remain constant for a small perturbation to the system). However, this will only work if there is a ring PZT at the back of the cavity to allow the laser to reflect off the back of the mirror. This will not always be possible and, even if it were, ideally we would use the same set-up for all measurements.

Another option is to measure the closed-loop transfer function, as shown in Fig. 4.2(a), where the input for the network analyser is taken from just before the controller, and the output is combined with the output of the controller. If only very small gain is used, especially on the I controller as it also affects the phase of the measurement, we can approximate $H \approx 0$ and therefore we obtain $T \approx G$ from Eq. 4.4. This can be seen by comparing Fig. 4.2(b) with Fig. 4.7(a), with traces of G measured using a Michelson method, and T using minimal controller gain, respectively (though for different cavities). Both traces show the same characteristic features of the plant: flat frequency response at low frequencies, for both magnitude and phase, before the magnitude increases towards the system resonance while the phase heads towards $-\pi$.

The reason for measuring G becomes apparent when the stability and bandwidth of the system need to be determined. Both of these depend on GH and, thankfully, H is easy to simulate. Figure 4.7(b) shows a comparison of H calculated from T and G and a version simulated from Eq. 4.9. Once G and T are measured they can be used to examine the lock. This will be discussed in the following section. One other advantage of measuring G is to help with calibration. In Sec. 4.2 it was discussed how, hopefully, the only ‘‘interesting’’ element of G (at least as far as frequency responses go) is the cavity itself G' and so we could write $G = \mathcal{K} \cdot G'$. This \mathcal{K} will lead to a constant offset on the Bode plot for G' , and there will be a similar effect for H and T . One way of negating this is to determine the value of these constants, which turns out to be quite easy if G has been measured. This is because, at low frequencies, $G' = 1$ and therefore \mathcal{K} can be determined by inspection from G bode plot at these frequencies. This value should be the same for T and should not change over time. The same is true for H , and the constant offset here can be determined by rearranging Eq. 4.4 to give

$$H(\omega) = \frac{G - T}{G \cdot T}. \quad (4.10)$$

This can then be compared to the simulated version and the extra gain inferred (see Fig. 4.7(b)). These normalising options are built into the program and were used for the following analysis.

The other method for characterising the lock response is RMS noise. This information is easily attained from the RT scope function for any lock.




Figure 4.8: Lock Optimisation Protocol. (a) Magnitude plot of system transfer function for an OPO cavity with controller gain increasing from (i) to (iv). (b) Normalised RMS values for these traces for both the error (ES) and dc (DC) signals. Error bars derived from standard deviation across ten runs used for each point. Also shown is the sequence in which the different gains are added, the bandwidth of the system (ω_B) and phase margin (PM) measured by the code.

4.4.2 Lock Optimisation

In our analysis we used the two tools mentioned above to investigate the individual locks: the system transfer function and the RMS of the error and dc signals, both measured internally using the code. For the transfer functions, the three key indicators were the amount of noise suppression at low frequencies, the range of frequencies for which noise was suppressed (i.e., bandwidth) and the stability of the lock as measured by the phase margin, as discussed in Sec. 4.2.

Figure 4.8 illustrates an optimisation procedure developed for a generic lock applied to the OPO:

1. **Measuring G .** This need only be undertaken once for each lock and then reloaded into the code later if necessary. As mentioned previously, to measure G we need to use as little gain as possible, especially with I gain. Normally, the minimal I gain needed to lock the system delivers the best approximations of G , with the option of adding P gain and decreasing I gain available if the output Bode plots do not display the same behaviour as Fig. 4.2(b) or Fig. 4.7(a). This step should also be used to determine the offset of T .
2. **I Gain.** As the system has no memory of its current position apart from where the scan stopped (which will be closer to the sides of the peak than the middle) I gain is important. This is because I gain has memory, meaning that if the value of the error signal is zero the I value will not change from its previous value, hopefully keeping the system at the centre of the peak. P gain, on the other hand, would be zero at the same point, moving the system away from the centre. I gain is also handy in that, if the correct amount of gain is used, it reduces noise at low frequencies while not exciting any system resonances. The amount of gain also determines the bandwidth of the system ($GH = 1$), therefore the I gain should be increased as much as possible until it either excites the system resonance or the phase margin drops below its desired value (normally taken to be 30°).
3. **P Gain.** Once I gain has been increased as much as possible it is time to add P gain. Normally this is simply added until it excites the resonance above a certain value or the system is seen to visibly oscillate. This P gain may make it possible to increase the I gain further, making Steps 3 and 4 an iterative process.
4. **I² Gain.** Once the values for I and P gains have been optimised, I² gain can be added to further reduce the system noise at low frequencies. It is important to be aware of any system instabilities that can be caused by increasing the gain due to the $-\pi$ phase delay.
5. **RMS Measurements.** To determine whether these increases in noise suppression are actually having an effect on the power fluctuation of the laser output noise measurements should be performed. As with the LabVIEW network analyser it is best to take a number of traces to make sure that the values are consistent. The sampling time for these RMS values (on the order of one second) depends on the sampling rate of the digital scope and therefore can only really be used for comparison purposes.

Figure 4.8(a) shows the system transfer functions for this process and the reduction in noise that occurs at low frequencies. Figure 4.8(b) shows the average RMS for these

traces taken over ten runs, illustrating the reduction in noise as more gain is added to the system. Also shown on this figure is the bandwidth and phase margin measured by the system. As can be seen, though the noise suppression measured on the transfer function traces does not increase significantly after the initial I gain is added, the bandwidth does increase while the RMS decreases with increasing gain. In all cases the phase margin was kept above $\pi/6$.

Though the amount and ratios of the gains required to optimise each individual lock will differ due to different plant resonances etc., the combination of the P, I, and I² controllers allow for flexibility when optimising.

4.4.3 Long-Term Stability and Comparison with Analog PI Controller

The long-term stability of an OPO cavity was investigated over a period of one hour. To accomplish this, both the error and dc signals for the OPO were measured at one second intervals using the code. The OPO was locked firstly with the internally programmed PII controlled and then with an external analog PI controller developed in-house. These analog PIs have been used previously on many quantum optics experiments, including the one described in Ref. [242]. The results of this comparison are shown in Figs 4.9(a) and (b) for the error and dc signals respectively, with the former converted to frequency and the latter normalised to the maximum peak height.

As can be seen, both the analog and digital controllers are stable over the period of one hour, though at one point, approximately 2200 s into the run, it can be seen that the analog system jumped slightly, but not enough to lose lock. This stability was despite the fact that during an hour the laser resonant frequency was found to move by approximately 175 MHz (measured by monitoring the PII controller output during the digital stability run). This is shown in Fig. 4.9(c).

The amount of noise suppression over this time was again measured using transfer functions and RMS values, with the former shown in Fig. 4.9(d). From the RMS values included in Figs 4.9(a)(i) and (ii) it can be seen that the RMS for the digital error signal trace was approximately 0.3 MHz, as opposed to the 0.4 MHz for the analog controller. For the dc RMS, even removing the perturbed region around 2200 s from the analog trace, the RMS was approximately the same (S_i being normalised to the total size of the peak and S_{dev} is taken to be the peak minima). Also, from Fig. 4.9(d) it can be seen that the digital controller has 50 dB greater noise suppression at low frequencies.

These results will vary depending on the different analog and digital controllers developed and the gains used. For instance, the extra noise suppression at lower frequencies for the digital controller can, in part, be attributed to the I² component, which was not included in its analog counterpart. However, it is much easier to alter a digital, rather than an analog, controller and this comparison verifies the efficacy and long-term stability of this general digital controller.

4.4.4 Other Optimisations

Apart from the locking parameters themselves there are many other ways to improve the overall quality of the lock. Some of these are briefly discussed below.

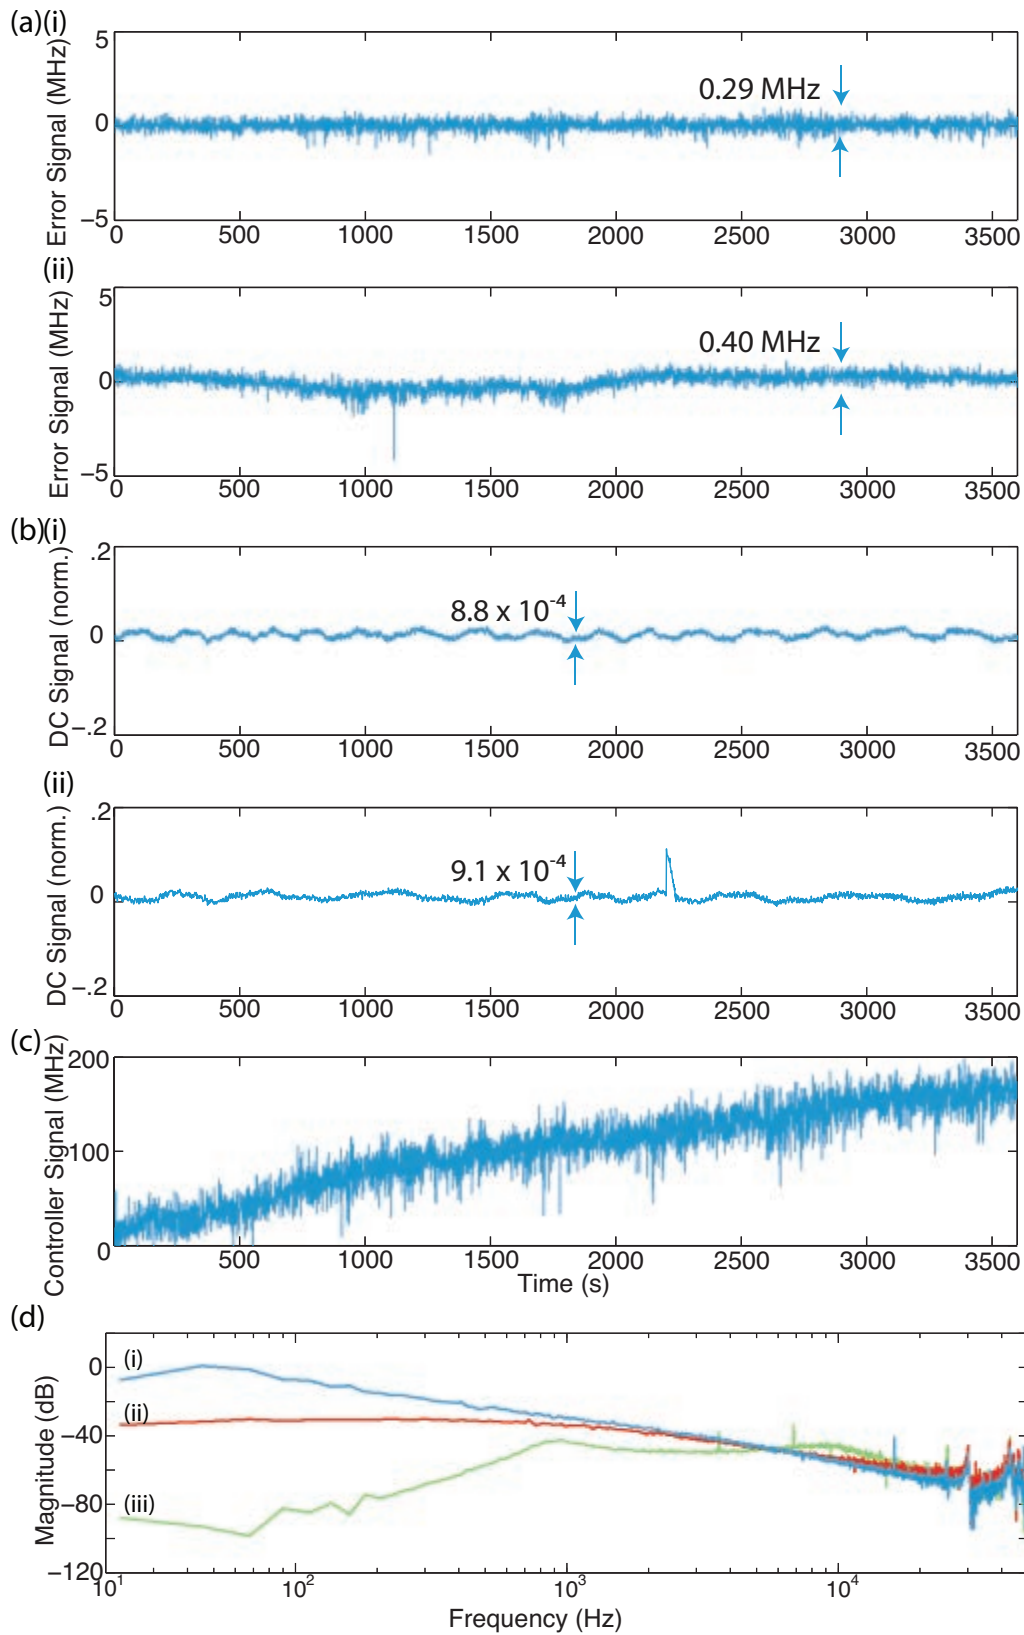


Figure 4.9: Analog and Digital Controller Comparison. (a) Error signal, measured in megahertz, and (b) dc signal, normalised to total peak height, for (i) a digital PII and (ii) analog PI controller, measured over a one hour period. (c) The digital PII controller signal measured during the digital run showing the frequency drift of the laser that was compensated for. (d) Magnitude plots of the system transfer function for (i) no locking, (ii) for locking with the analog PI controller, and (iii) for locking with the digital PII controller.

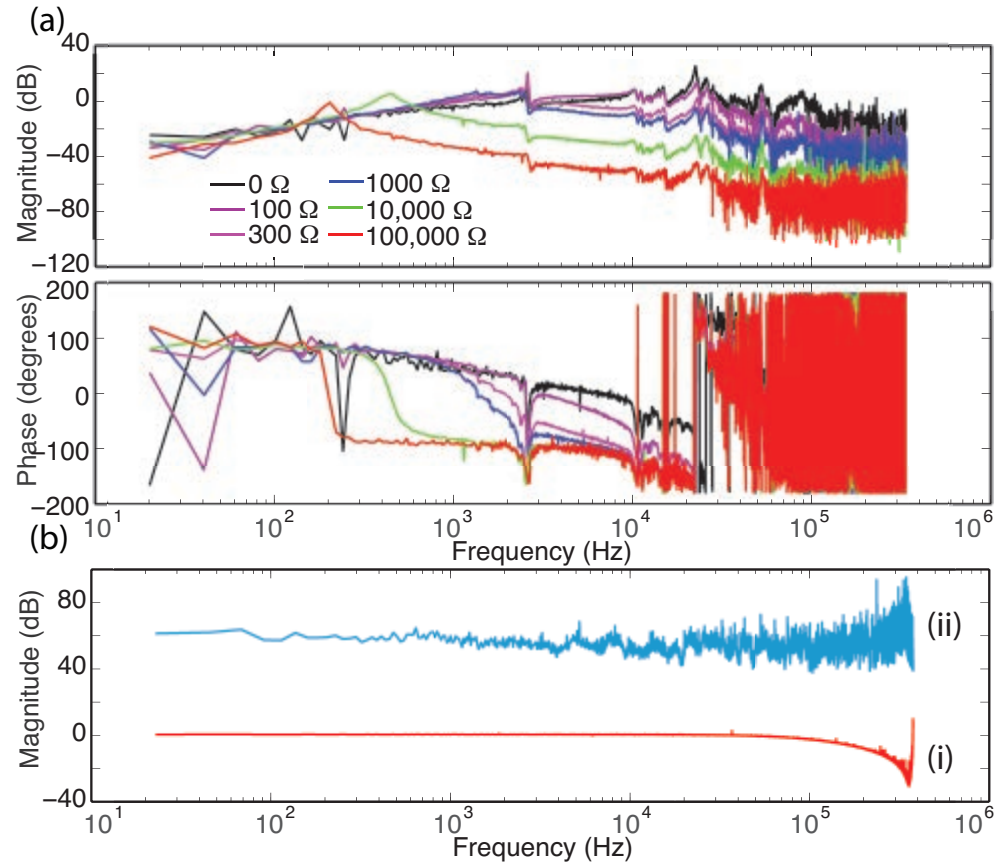


Figure 4.10: Low-Pass Filtering. (a) Bode plot of T for a ring cavity with different resistors attached after the HV amplifier (see key). (b) Magnitude-only plot of infinite-impulse-response (IIR) filter created in LabVIEW for two frequencies, (i) 124.7 kHz and (ii) 124 kHz, with a cut-off frequency of 124.5 kHz.

Low-Pass Filters and Resistors

One way to reduce the ringing at higher frequencies associated with increasing P gain is to include a low-pass filter (LPF) between the controller and the PZT. As a low-pass filter consists of a resistor and capacitor in series, and as a typical PZT has an inherent capacitance on the order of 1 nF, adding a resistor can also be approximated to a LPF, with the cut-off frequency f_c being given by [249]

$$f_c = \frac{1}{2\pi CR}, \quad (4.11)$$

where C is the capacitance of the system and R the resistance. Figure 4.10(a) shows the effect of adding different resistors between the HV amplifier and the PZT. As can be seen, the higher the resistance, the lower the magnitude of T at frequencies higher than f_c . However the effective capacitance seems to decrease with increasing R (from approximately 10^{-7} to 10^{-9} F for the range of resistances shown) and therefore adding a resistor after the HV amplifier must also be having some other effect. For instance, increasing R decreases the voltage supplied to the PZT and, therefore, the system gain. It will also increase the controllers sensitivity by decreasing the range of the HV amplifier

(see following section).

Using a LPF (with capacitance on the order of a microfarad) is a better alternative as it allows for more exact control of f_c and will not alter controller gain below the cut-off frequency. The choice of an optimal value for f_c (by choosing the correct combination of R and C) is an interesting one and depends on what are the main characteristics desired for the lock. For instance, if reducing noise at low frequencies (less than a kilohertz) is important, then a filter with $f_c \approx 10$ Hz is a good option. If this is not a priority, choosing f_c anywhere up to the resonant frequency will help the lock by allowing for maximum gain at the lower frequencies. One side-effect of using a LPF is that it can interfere with the autolock mechanism, as the abrupt change from scanning to locking functionality can be interfered with if f_c is too low (or, for the case of using a resistor, R is too high). This issue can be overcome by reducing the scan speed.

An alternative to a physical LPF is to use a digital one programmed in LabVIEW. Again this is made easy as LabVIEW contains an inbuilt tool to determine parameters for infinite-impulse-response (IIR)-type filters. However, these parameters only work for $f_c > 0.166f_s$, as shown in Fig. 4.10(b), which for the code described above would be approximately 124.5 kHz. This is not a useful cut-off frequency for the systems of interest to us and therefore this digital filtering method was not used. If a digital filter was desired, again due to its flexibility and ease of use, the code would have to be run at one quarter of its current speed to obtain a “useful” $f_c \approx 30$ kHz, though care must be taken with this approach as reducing f_s would increase the phase delay at higher frequencies.

Digital Noise

Unfortunately there are some disadvantages to go along with the many advantages of digital locking, and one of these is digitisation noise. This is due to the finite sensitivity of the system for both the input and output signals and may introduce additional noise into the system. In the case of the FPGA, 2^{16} bits are used for both signals, each with a voltage range of ± 10 V. Each bit therefore represents $20/2^{16} = 3.05 \cdot 10^{-4}$ V. For a “typical” error signal entering the FPGA, the peak-to-peak range is on the order of $V_{pp} = 0.5$ -1 V, meaning that the percentage error caused by the digitisation is $3.05 \cdot 10^{-4}/1 \cdot 100 = 0.03\%$. This does not seem very large, though it can be further reduced by increasing the size of the error signal to the maximum possible ($V_{pp} = 20$ V). The difference in dc RMS values for error signals ranging in size from 0.3 to 18 V are shown in Fig. 4.11(a). To keep the actual controller gain ($k_H \cdot V_{pp}$) constant, whenever the error signal height was changed (by approximately a factor of two each time) k_H was doubled/halved as appropriate. As can be seen from Fig. 4.11(a), there is a negligible effect within the error bars (though a trend towards higher RMS can be seen as $V_{pp} \rightarrow 0$) and therefore this effect can be ignored except for the most sensitive of applications.

Similar to the input scenario, the output controller signal should also span the full ± 10 V to minimise digitisation noise. However, we also want the scanning range of the signal to be at least one free spectral range (FSR). These two features cannot both be achieved simultaneously through one output of the FPGA. If we satisfy the second but not the first, then the sensitivity of the locking will be limited as, if the finesse of the cavity is approximately 1000, the peak full-width-half-maximum (FWHM) will only be $2^{16}/1000 \approx 65$ points (see Fig. 4.11(b)). One solution to this problem is to use *two* outputs for each lock: one for the scan and one for the lock. The idea here is to set the

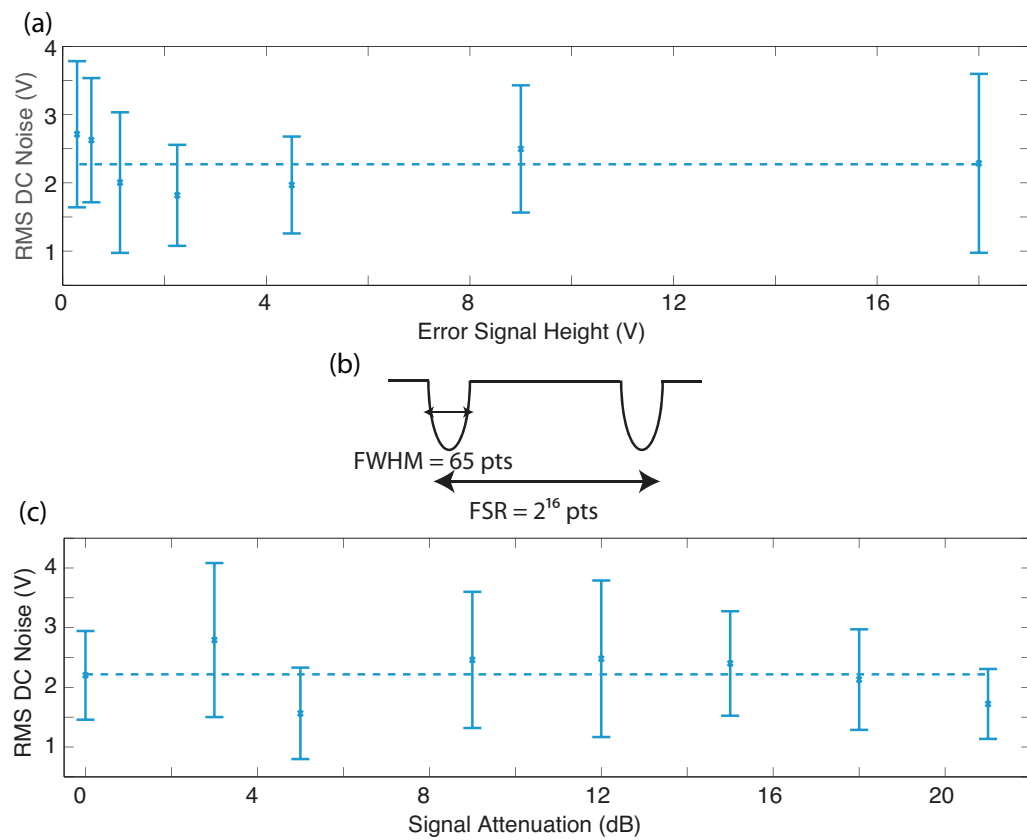


Figure 4.11: Digital Noise. (a) RMS measurements of the dc signal for variously sized error signals. (b) Schematic of the dc signal as a function of cavity position showing the width, in points, of both an FSR and the FWHM of one peak (not to scale). (c) RMS measurements of the dc signal for varying attenuation applied after the LabVIEW output and before the HV amplifier.

scan range to be one, or better two, FSRs (ensuring that at all times at least one peak is not on the outskirts of the scan range), and using analog attenuation for the lock output to increase the sensitivity of the locking. This method works well in principle with the scan and lock system: the scan output, with its large range, can always find the peak; once this happens it will set the dc offset around which the system will lock with the attenuated lock signal; and, if long term drift of the system proves too much for the lock signal to negate over a long period of time (i.e., hours), then the system will simply relock and start again. In this case (assuming a scan range of three FSR) the number of points p for the FWHM will be given by

$$p = 2^{16}/2000 \cdot 10^{att/10} \approx 33 \cdot 10^{att/10}, \quad (4.12)$$

where att is the analog attenuation applied to the locking signal in dB. The effect on the RMS of the signal with varying amounts of attenuation is shown in Fig. 4.11(c) (again with increased k_H to compensate for loss in controller gain due to the attenuation) and, as can be seen, this again has a negligible effect on the locking quality within the error bars. A slight trend towards lower RMS values can be seen as att increased, however, more data points would be needed to determine if this is a real effect and the system refused to lock with $att > 21$ dB. This means that the two output method is unnecessary for most applications, leaving open the possibility of using the remaining output for other functions such as a trigger signal for an external scope, etc.

4.5 Quantum Measurements

So far we have presented results which are not specific to the field of quantum optics. In this section, a bow-tie OPO cavity was inserted into a larger quantum optics experiment designed to characterise optical Schrödinger cat states (see, for instance, Ref. [250]). This experiment, described in detail in Ref. [251] and shown in Fig. 4.12(a), consisted of two mode-cleaner cavities (one at 532 nm and one at 1064 nm), the OPO - which was used to produced squeezing on the 1064 nm beam, two phase locks, a double homodyne, and a tomographic homodyne lock to allow locking to any quadrature value. This was all controlled using the hardware and software described in Sec. 4.3 and integrated into one user-controlled RT code.

As in many other quantum optics experiments, it was desirable to only take data when all components of the system were locked. Also, in this case, vast quantities of data (hundreds of gigabits) were required to be able to produce the necessary tomographic reconstructions of the optical Schrödinger cats, requiring many hours of acquisition. It was also necessary to compare the squeezing data traces with dark noise (DN - signal and local oscillator blocked) and shot noise (SN - only signal blocked) traces at various times during data collection for normalisation purposes.

To accommodate these requirements, a second RT code was developed to take data using two data acquisition cards (NI PXI-5124). The data acquisition code communicated with the locking code to determine whether all components were locked, as well as to change locking parameters as desired. For instance, the phase of the single homodyne lock had to be changed to allow for the tomographic reconstruction, and blocking of the various fields present using motors controlled via the locking FPGA was necessary for collecting DN and SN traces. The data was saved to the hard drive of the PXI controller and later extracted for examination.

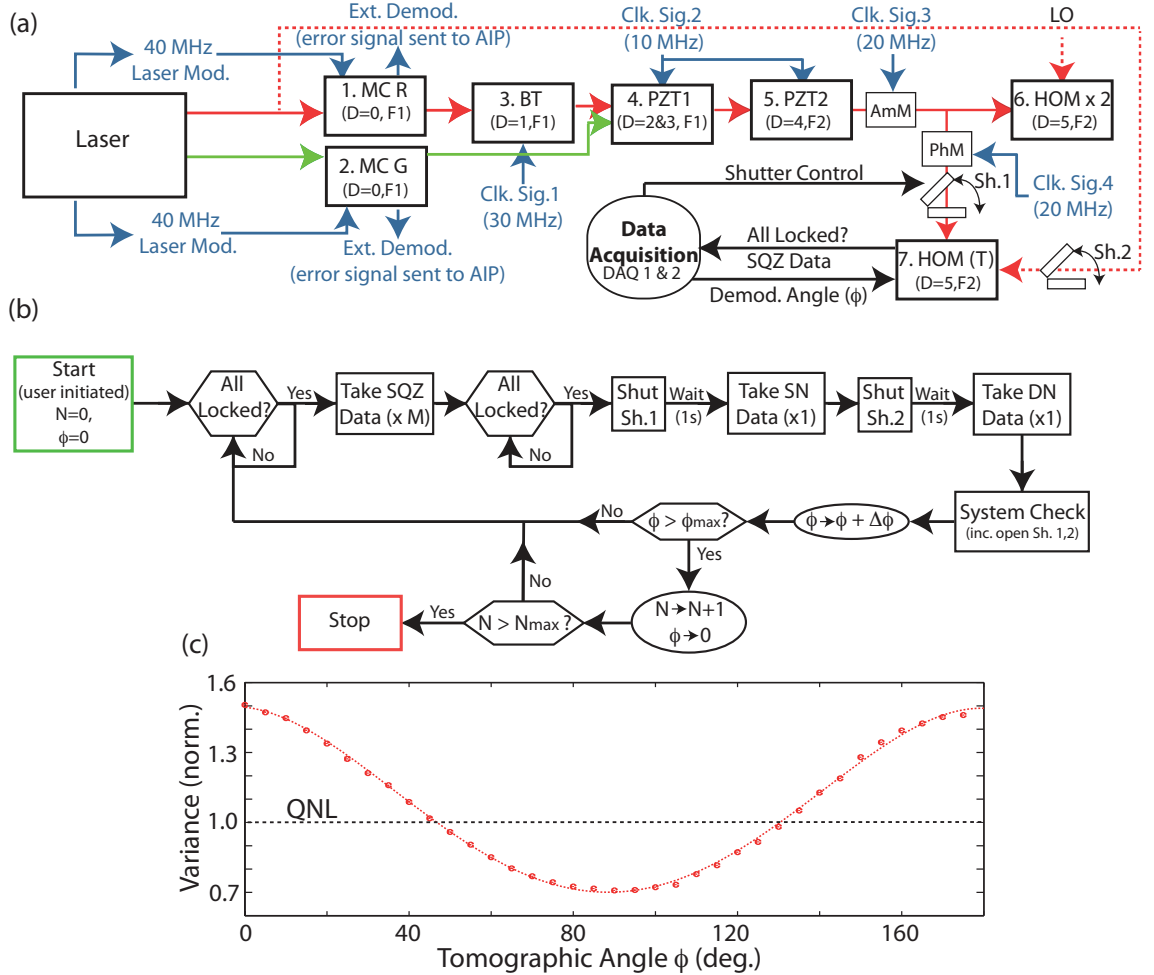


Figure 4.12: Schrödinger Cat Experiment. (a) Schematic of cavity set-up with: MC - mode-cleaner cavity (R - red, G - green); BT - bow-tie cavity; PZT - piezo-transducer phase lock; HOM - homodyne lock (T - tomographic, x 2 - double homodyne); AmM - amplitude modulator; PhM - phase modulator; Sh - mechanical shutter; SQZ Data - squeezing data; D - lock dependencies (0 = no dependencies); F - FPGA lock is controlled from; Ext. Demod. - external demodulation carried out for these locks to produce the error signal; AIP - analog input; DAQ - data acquisition card; Clk. Sig. - clock signal; and LO - local oscillator. (b) Flow diagram of data acquisition procedure used to take squeezing data (SQZ) for various angles of the tomographic lock ϕ , as well as take dark noise (DN) and shot noise (SN) traces. Values for the number of experimental runs N , number of squeezing traces for each tomographic angle M , and change in tomographic lock angle between runs $\Delta\phi$ are defined by the user. (c) Plot of variance in noise normalised to the quantum noise limit (QNL), as a function of tomographic lock angle taken using the above procedure with a total of 1.2 million points per angle, and $\Delta\phi = \pi/36$. Circles correspond to data points, the dashed line corresponds to a theoretical fit.

A flow diagram of the data collection protocol is shown in Fig. 4.12(b). Here the user determined the number of data traces to take for each tomographic lock angle M , as well as the change in tomographic lock angle between measurements $\Delta\phi$, the maximum tomographic lock angle ϕ_{max} , and the number of experimental runs N for $\phi = 0 \rightarrow \phi_{max}$. It should be noted that the waiting time programmed into the code was to allow time for the shutters to close or open, and the system check included such functions as resetting error signal offsets so that all locks remained stable during the hours of data collection. If a lock did fail during the measurement of one angle, the protocol would reset and was attempted again for that angle until it was successful. Figure 4.12(c) shows an example of squeezing data collected from one such run with $\Delta\phi = \pi/36$. Approximately 1.5 dB of noise suppression below the quantum noise limit (QNL) could be seen for a tomographic angle of $\phi = 90^\circ$, accompanied by an increase of 1.8 dB in the opposite quadrature.

This implementation of a digital control algorithm on a uniquely quantum system highlights the benefits digital control can bring to quantum optics experiments: the integration of information allowed for data to only be taken when all system components were locked to ensure only useful data was collected; the automation of the locks meant that the experiment could be left unattended for many hours to collect the required amount of data as quickly as possible; and the added flexibility allowed for automated acquisition of SN and DN, as well as the ability to alter the homodyne angle to perform the tomographic measurements.

4.6 Conclusions

In conclusion, we have presented a digital locking system that allows for automatic and sequential locking and is easily scalable. This code was programmed using LabVIEW software and is free to download. We have shown how the inbuilt locking analysis tools can be used to help optimise individual locks, analysed the code and its behaviour on multiple types of cavities in great detail, and demonstrated the efficacy and long term stability of the digital controller. Finally, we used an example of a experiment used to characterise optical Schrödinger cat states to illustrate the advantages of digital control for quantum optics experiments.

Part II

Warm Vapour Experiments

Introduction to Light Storage with Warm Atomic Ensembles

Daffy Duck: “Say, now you’re cooking with gas.”
 - *The Wise Quacking Duck*, Warner Bros., May 1943, Movie.

5.1 Introduction - Why Warm Atoms are Hot Stuff

In the following chapters of this thesis we will be using warm atomic ensembles to test various properties of the Λ -GEM scheme, including its ability to spectrally manipulate information and store images. Warm atomic ensembles, usually consisting of alkali atoms in vapour cells (see Fig. 5.1), are an ideal test-bed for quantum memory protocols for a number of reasons. Firstly, they are relatively cheap and easy to use - not requiring the ultra-high vacuum systems needed for cold atom or ion-based quantum memories, or the cryostats that are needed for most solid-state ensembles. This, in turn, makes them robust, a property that will need to be considered when trying to engineer actual quantum networks containing quantum memories in remote locations. Importantly for high efficiency applications they can have ultra-high optical depths (ODs) in the thousands [7, 252]. There is also the possibility that in the future vapour cells could be miniaturised [253], put on a chip [254], or the vapour trapped in a hollowcore wave-guide, where ODs of $3 \cdot 10^4$ have been measured for Cs [255].

This chapter presents a literature review of some of the standout results achieved using warm-atom-based quantum memory protocols in Sec. 5.2 and then presents some of the experimental issues faced when using them in Sec. 5.3.




Figure 5.1: Warm Vapour Cell. Natural mixture Rb vapour cell, heated to approximately 150°C and illuminated with approximately 300 mW of light at 795 nm. (Photo credit: M. Hosseini and B. C. Buchler).

5.2 Literature Review of Warm Atom Light Storage Experiments: 2001-Present

Starting at the beginning, two near-simultaneous publications emerged in January 2001 demonstrating light storage for the first time with electromagnetically-induced transparency (EIT). One of these used a cold atomic ensemble [121] (see Sec. 9.4.1) and one was carried out using a 4-cm-long vapour cell of ^{87}Rb atoms and 5 Torr of He buffer gas¹ heated between 70-90°C. This proof-of-principle work, by D. F. Phillips et al., showed storage for 200 μs of approximately 10-30- μs -long pulses, with an efficiency on the order of 10% [122].

Since then, many warm-atom-based EIT experiments have been carried out. Some of note include the generation and then storage of single-photons with two warm vapour cells by M. D. Eisaman et al. in 2005. Here the authors generated a single-photon with a DLCZ scheme (see Sec. 1.5.4) in one 4.5-cm-long ^{87}Rb vapour cell with 7 Torr of Ne buffer gas, and then stored it in another 4.5-cm-long ^{87}Rb vapour cell, this time with 8 Torr of Ne buffer gas [126]. The EIT storage had an exponential decay with a coherence time (i.e., $1/e$ decay time) of 1 μs , and a maximum efficiency of 10% for the approximately 100-ns-long pulses. There was also the work by J. Appel et al. who, in 2008, demonstrated storage of a 1.86 dB squeezed-vacuum pulse, with up to 0.21 dB retrieved [129]. This work was carried out in a ^{87}Rb cell with 10 Torr of Ne buffer gas heated to 65°C and, with this medium, they achieved a coherence time of 1.3 μs . Finally, work by N. B. Phillips et al. in 2008 extended the efficiency for EIT in warm vapours [124]. To do this, the group used a 7.5-cm-long ^{87}Rb cell with 30 Torr of Ne buffer gas, heated between 45-75°C. Then, using approximately 10- μs -long pulses, they implemented iterative optimisation algorithms [256, 257] to reach a maximum efficiency of 43% with a storage time of 100 μs . The coherence time of their system was 500 μs . This efficiency record was only broken in 2011 with cold atoms (see Sec. 9.4.1).

Returning to the DLCZ scheme mentioned in the previous paragraph, the first proof-of-principle storage of anti-Stokes photons was achieved in 2003 using warm atoms. C. H. van der Wal et al. used a 4-cm-long ^{87}Rb vapour cell heated to 85°C, where a write beam was used to create Stokes photons before waiting and then using a read beam to produce the correlated anti-Stokes photons [145]. The authors were able to show a maximum efficiency for the anti-Stokes recall of 30% and a coherence time of 1 μs .

Another type of memory that has, to date, been implemented exclusively in warm atoms is the Raman memory developed by the group of I. A. Walmsley in Oxford. The first demonstration of a Raman memory was achieved using a 7-cm-long Cs cell with 20 Torr Ne buffer gas heated to 62.5°C by K. F. Reim et al. in 2010 [7]. With this system they were initially able to achieve 15% efficiency for light stored up to 12 ns, but with an impressive bandwidth of 1.5 GHz (i.e., pulse-width down to 300 ps). This was improved the following year with single-photon-level storage demonstrated with an efficiency up to 30% and a coherence time of 1.5 μs , to give a delay-bandwidth product of 2,500 [142].

We will discuss the detrimental effects of four-wave mixing (FWM) later in this chapter. However, FWM has also been used to delay and store information in warm vapours. For instance, in 2007 V. Boyer et al. were able to demonstrate a delay of up to 40 ns for a

¹More on the role of buffer gas in Sec. 5.3.3.




Figure 5.2: High Efficiency and Multi-Pulse Λ -GEM Storage. (i) Inputs and (ii) echoes for (a) high efficiency storage (see labels for efficiency) and (b) multi-pulse storage (note $\times 20$ label), dashed line indicates magnetic field switching point. Figure adapted from Ref. [4].

70-ns-long pulse without much distortion inside a 2.5-cm-long ^{85}Rb cell heated between 90-140°C [193]. In 2008, the same group was able to show the creation of entangled images, again in warm ^{85}Rb [189] and, in 2009, were then able to tuneably delay these images for up to 32 ns in a secondary ^{85}Rb warm vapour cell [194]. In terms of storage of light using FWM, R. M. Camacho et al. were able to store pulses using a mixture of EIT and FWM in a natural-mixture Rb cell with 20 Torr Ne buffer gas heated to 70°C [195]. They were able to achieve efficiencies of a few percent, and a coherence time of 120 μs . FWM in warm vapour cells has also led to the production of microsecond-length entangled pulses [190] and squeezing of nanosecond-length pulses [191, 192].

Of course, many other protocols have also used warm atoms for light storage. For instance, measurement and feedback was used to map light into two Cs warm vapour cells by B. Julsgaard et al. [196] and a memory for continuous-variable squeezed states was created with another warm Cs cell and feedback by K. Jensen et al. [197].

We will now finish our discussion of warm atom quantum memories by discussing the Λ -GEM scheme. Λ -GEM was first demonstrated in 2008 using two m_F states of the $F = 2$ level in a 7.5-cm-long ^{87}Rb cell with 5 Torr He buffer gas, heated to approximately 80°C. This set-up achieved efficiencies of only a few percent, with a coherence time of microseconds and a delay-bandwidth product of four [183, 182]. Since then, by changing the polarisations and frequencies of the probe and coupling fields (for more information see Ref. [258]) as well as the cell, to a 20-cm-long cell with 0.5 Torr Kr buffer gas, efficiencies have been increased up to 87% [4]. This is, to date, the highest efficiency for an unconditional memory. Coherence times have been improved up to 47 μs achieved (see Sec. 8.4), and up to 20-pulse storage has also been demonstrated. The memory efficiency and mutli-pulse storage results are shown in Fig. 5.2.

5.3 Not So Hot - Decoherence Effects in Warm Vapours

Though from the above results it can be seen how useful warm vapour cells are for quantum memory applications, there are also draw backs to using a system where atoms are moving at hundreds of meters per second. In this section we will discuss some of the major effects

limiting quantum memories, focusing but not limiting ourselves to warm atomic systems.

5.3.1 Doppler Broadening

When atoms are moving at a non-negligible speed in relation to light fields, they will experience a change in the laser frequency they “see” in their rest frame. This is called the Doppler shift, following [204]

$$\omega_a = \omega_l - \mathbf{k}_l \cdot \mathbf{v}_a, \quad (5.1)$$

where ω_l is the frequency of the laser in the reference frame of the lab, \mathbf{k}_l is the wave-vector of the laser field ($\mathbf{k}_l = (2\pi/\lambda_l)\mathbf{r}$, where λ_l is the wavelength of the laser and \mathbf{r} is the direction vector), \mathbf{v}_a is the velocity of the atom and ω_a is the frequency of the laser in the reference frame of the atom. The velocity distribution of the atoms in a vapour cell can be determined by knowing the temperature of the cell and using the Boltzmann distribution. The most probable velocity v_{pr} for a given temperature being [204]

$$v_{pr} = \sqrt{\frac{2k_B T}{M}} \approx 260 \text{ m s}^{-1}, \quad (5.2)$$

where k_B is the Boltzmann constant, T is the temperature in kelvin taken here to be 80°C, and M is the mass of ^{87}Rb . Here we will approximate the Boltzmann distribution to a Lorentzian, so that we can determine the full-width-half-maximum (FWHM) of the Doppler-broadened profile $2\Gamma_D$. This approximation is allowable as the only real difference between the two distributions is on the wings, where there are not many atoms present [259]. In this case, assuming a ^{87}Rb sample at 80°C with no buffer gas and a laser wavelength of 795 nm, we can use the following equation to determine the Doppler width of our memory [139]

$$\Gamma_{D,\Delta} = \frac{2\pi}{\lambda_l} \sqrt{\frac{2\ln[2]k_B T}{M}} \approx 2\pi \cdot 160 \text{ MHz}. \quad (5.3)$$

If we were only using a two-level atom then this would be disastrous for storage times (as coherence time $t_c \propto 1/\Gamma$). However, Λ -GEM uses two fields: a weak probe and strong coupling field, with wave-vectors \mathbf{k}_p and \mathbf{k}_c respectively, as is the case with EIT and Raman memories. In this case, it is the *difference* in frequency between the two fields seen by the atoms that will determine the change in two-photon resonance δ via

$$\Delta\delta = (\mathbf{k}_p - \mathbf{k}_c) \cdot \mathbf{v}_a. \quad (5.4)$$

If we assume that the probe and coupling field are co-propagating, or only at a very slight angle, are on the two hyperfine level $F = 1$ and $F = 2$ respectively with $\Delta_{hfs} = 6.8$ GHz, and the vapour cell has a Doppler width $\Gamma_D \gg \gamma$ (the natural linewidth $\gamma \approx 2\pi \cdot 6$ MHz for ^{87}Rb), then we can determine the effect of the Doppler shift on the width of the quasi-two-level transition to be [260]

$$\Gamma_{D,\delta} = \frac{\gamma}{\omega_l} \Delta_{hfs} \approx 2\pi \cdot 100 \text{ Hz}, \quad (5.5)$$

in the high temperature limit. If this were the only factor limiting storage in a vapour cell, then storage times on the order of milliseconds should be achievable without any further

work. However, we must still ensure that we have a one-photon detuning $\Delta \gg \Gamma_{D,\Delta}$ for our Raman system to avoid coupling light to the excited state, which will lead to scattering effects (see Sec. 5.3.4).

5.3.2 Diffusion

Sadly, unlike solid-state systems or optical lattices, with non-negligible velocities we will also have atoms moving into and out of the interaction region. The time-of-flight of the atoms in the interaction volume in many cases limits the coherence time of the system. For instance, from the experiments described in the previous section, Ref.s [126] and [145] both saw coherence times limited to the order of one microsecond due to diffusion effects. We can do a rough calculation of the time-of-flight of an atom inside a beam with a radius of 1.5 mm (as will be used in Ch. 8) then, using v_{pr} , we find that an atom will spend on average about 6 μs inside the beam using the scenario shown in Fig. 5.3(a) (i.e., assuming no interatomic collisions).

To help quantify the effect of diffusion in dispersing atomic excitations, we can define a diffusion constant D that determines the additional decoherence due to the movement of atoms by [261]

$$\frac{\partial}{\partial t} \hat{\sigma}_{jk} = D \nabla^2 \hat{\sigma}_{jk}, \quad (5.6)$$

where $\hat{\sigma}_{jk}$, $j \neq k$ are the atomic coherence components of the ensemble density matrix and for $j = k$ they are the populations. These terms would add to the time-evolution derived for the coherences and populations in Sec. 2.5.3. Assuming a larger coupling beam waist than probe beam, diffusion will also lead to expansion of a stored probe beam. The retrieved probe waist after time t , $w_p(t)$, being determined by [262]

$$w_p(t) = \sqrt{w_p^2(0) + 4Dt}, \quad (5.7)$$

where $w_p(0)$ is the initial probe waist. For warm-atom EIT systems investigating image storage, D has been found to be approximately $10 \text{ cm}^2 \text{ s}^{-1}$ [262, 263]. D has been investigated for Λ -GEM with 0.5 Torr Kr buffer gas, and was found to be between 13 and $65 \text{ cm}^2 \text{ s}^{-1}$ without the coupling field on during storage [264] (more on this in Sec. 5.3.4).

There is also a counter-intuitive effect associated with atoms moving into and out of the interaction zone. It has been demonstrated that the motion of atoms out of, and then back into, the beam can actually lead to a narrowing of the Doppler-broadened transmission line in EIT [265, 266] as opposed to the broadening we would expect (see Eq. 3.22). This Ramsey narrowing has led to a 40-fold decrease in EIT peak size, from 26 kHz to 600 Hz in Ref. [265], with a final coherence time of 500 μs . The width of the narrowed peak is still limited by other motion-induced effects such as Doppler broadening (previous section) and collisions (next section) and, therefore, to achieve storage times approaching seconds we will still need to reduce atomic motion as much as possible.

So far we have only discussed transverse diffusion. However, due to the longitudinal gradient in GEM, we will also have to consider longitudinal diffusion as, if atoms move up and down the memory, they will rephase at different times as the phase matching conditions will not be perfect, leading to frequency blurring and therefore inefficient recall. To calculate the magnitude of this effect we can look at the change in frequency experienced




Figure 5.3: Diffusion and Buffer Gas. (a) Diffusion of Rb atom with most probable velocity v_p out of the interaction defined by the probe waist W_p . (b) Effect of buffer gas, with many elastic collisions with Kr keeping the Rb atom in the interaction zone much longer, with mean free path \mathcal{L} .

by an atom if it moves a distance $2\sqrt{D}t_s$, known as the diffusion length,

$$\Delta\omega_{diff} = 2\pi \cdot 2\sqrt{D}t_s \cdot \eta \approx 2\pi \cdot 8 \text{ kHz}, \quad (5.8)$$

where we have taken the gradient η to be on the order of 50 kHz cm^{-1} , used the larger value of $D = 65 \text{ cm}^2 \text{ s}^{-1}$, and assumed a storage time of $t_s = 100 \mu\text{s}$.

We can compare this change in frequency to the bandwidth of the pulse, which for a 1 MHz total bandwidth will be approximately 1%. We can also compare the diffusion length to the wavelength of the spinwave ($1/\eta t_s$). For a storage time of $100 \mu\text{s}$ the ratio approaches 1:1 and therefore this effect will start to impact the storage around this time. If the gradient is turned off during storage, the spinwave will stop evolving in k -space (as illustrated using the polariton description of GEM in Sec. 3.4.3). For instance, in the situation described above, if we were to store a pulse with a total width of $5 \mu\text{s}$ and then turn off the gradient immediately after the pulse had entered the memory, the ratio of diffusion length to the spinwave wavelength would be 1:20.

5.3.3 Buffer Gas, Collisional Broadening and Anti-Relaxation Coating

Rather than reducing diffusion by cooling the atoms (and therefore lowering the OD), we could insert an inert gas into the vapour cell. The inert “buffer” gas will lead to elastic collisions with the alkali element and therefore preserve the atomic coherence while reducing the mean free path \mathcal{L} of the alkali element. This is shown in Fig. 5.3(b), with up to ten million collisions with the buffer gas possible before coherence is lost [267].

The time-of-flight broadening is inversely proportional to the buffer gas pressure, and has led to very narrow EIT linewidths being measured. For instance, S. Brandt et al. demonstrated a 42 Hz linewidth using a Cs vapour cell with 15 Torr of Ne. Similarly, M. Erhard et al. showed a 28 Hz linewidth for a ^{85}Rb vapour cell with 45 Torr of Ne. In terms of storage experiments, the EIT experiment by N. B. Phillips et al. described above used 30 Torr of Ne in their ^{87}Rb vapour cell to obtain their coherence time of $500 \mu\text{s}$.

However, the inclusion of buffer gas does come at a cost, and that cost is collisional broadening - where the excited state population can be shuffled in the presence of the light fields - following [267, 268]

$$\Gamma_{col,\Delta} = \frac{P\sigma_k\bar{v}}{k_B T}, \quad (5.9)$$

where P is the pressure of the buffer gas, σ_k is the kinetic cross-section, and \bar{v} is the average velocity of the atoms. This eventually leads to a trade-off between the two effects, as seen in the curve of linewidth vs buffer gas pressure in Ref. [269]. The one-photon broadening due to buffer gas for Rb-Kr collisions is $\Gamma_{col,\Delta} = 17.1 \text{ MHz Torr}^{-1}$ [270] at room temperature, scaling approximately as \sqrt{T} . This effect is similar to the Doppler broadening, where it is in the two-photon broadening that is important. However, in this case, as $\Gamma_{col,\Delta} \ll \Gamma_{D,\Delta}$, we can neglect the effect of $\Gamma_{col,\delta}$.

We can calculate \mathcal{L} using the collision rate and atomic speed via

$$\mathcal{L} = \frac{v_p}{\Gamma_{col,\Delta}} \approx 10 \text{ } \mu\text{m}. \quad (5.10)$$

The equation that governs the decay due to diffusion in the presence of buffer gas is then given by [271]

$$\Gamma_{diff} = 2.405^2 \frac{D}{w_p^2 + 6.8\mathcal{L}w_p} \approx 2\pi \cdot 3 \text{ kHz}, \quad (5.11)$$

where we have taken $D = 65 \text{ cm}^2 \text{ s}^{-1}$ for the Λ -GEM system (i.e., 80°C and 0.5 Torr Kr buffer gas) and $w_p \approx 1.5 \text{ mm}$ is the probe waist.

One final consideration is to do with the behaviour of atoms when they hit the cell wall and then return to the interaction region. If the atoms hit naked glass the collision will most probably be inelastic, leading to a loss of coherence. However, if the glass is coated with an anti-relaxation coating, then the coherence may be preserved and storage times lengthened. These coatings include paraffin [272] where, for example, a 100-fold increase in coherence times were seen for DLCZ-type experiment using warm ^{87}Rb [273], as well as Octadecyltrichlorosine [274], 1-nonadecene [275] or other alkene coatings [276]. With the last, alkene, coating a coherence time in excess of a minute was demonstrated, with the coating able to take 10^6 collisions, 100 times better than paraffin coating.

5.3.4 Power Broadening and Spontaneous Raman Scattering

An effect not exclusive to warm atomic systems, but worth paying attention to nevertheless, is power broadening. This effect occurs when an atomic resonance saturates with large numbers of photons, leading to reduced absorption [207]. The one-photon on-resonance power-broadened width will be of the form [277]

$$\Gamma_{P,\Delta} = \gamma \sqrt{1 + 2\Omega_c^2/\gamma^2} \approx 2\pi \cdot 42 \text{ kHz}, \quad (5.12)$$

where $\Omega_c = 2\pi \cdot 30 \text{ MHz}$ is the Rabi frequency of the coupling field. It is again important that this value remains small compared to the one-photon detuning. The two-photon power broadening is more complicated, as it must be calculated for the coupling field interacting with both the $F = 1$ and $F = 2$ levels. This is given by [268]

$$\Gamma_{P,\delta} = \sum_F \frac{\Omega_{c,F}^2}{\gamma + \Delta_F^2/\gamma} \approx 2\pi \cdot 3.5 \text{ kHz}, \quad (5.13)$$

where $\Delta_2 = 2\pi \cdot 2 \text{ GHz}$ (the one-photon detuning shown in Fig. 5.4(b)) and $\Delta_1 = -2\pi \cdot 4.8 \text{ GHz}$, with $\Omega_{c,1} = 2\pi \cdot 90 \text{ MHz}$ (due to Clebsch-Gordan coefficients - see Ref. [184]).

This is known as power broadening because $\Gamma_P \propto \Omega_c^2 \propto I_c$, the intensity of the coupling

field. It is interesting to note here that, for $\Delta \gg \gamma$, $\Gamma_{P,\delta} \propto (\Omega_c/\Delta)^2$. This is the same scaling as the Raman OD (see Eq. 3.54). This behaviour of the power broadening is one reason that, unless $\Omega_c/\Delta \ll 1$, increasing the coupling field power will not greatly improve the memory efficiency.

The increased decay rate due to the coupling field can also be explained in terms of spontaneous Raman scattering (SRS). With the large coupling field powers required for Λ -GEM in warm vapours (i.e., hundreds of milliwatts), and the large Doppler broadening of the upper level, there is a high probability that photons from the coupling field will be absorbed and then reemitted by the atoms. This is known as scattering, and will be discussed in detail in Sec. 9.3.3. There are two main components to the scattering: elastic or Rayleigh scattering, where atoms return to their initial state; and inelastic or Raman scattering where they do not. It is the latter that will have a larger effect on the coherence. SRS has been seen to increase noise in other quantum memory protocols implemented with a strong coupling field in warm vapour cells, as the Stokes/anti-Stokes photons generated are detected along with the probe [142, 278, 279].

After this analysis we can see that turning off the coupling field during Λ -GEM storage should improve the coherence time of the memory. This increase has been demonstrated experimentally to be nearly an order of magnitude [4, 252]. As the coupling field will, in most cases, be Gaussian and therefore not uniform across the atomic ensemble, the scattering will also lead to an increase in the transverse blurring of the probe if it is left on during storage as $\Gamma_P \propto I(\mathbf{r}, \mathbf{z})$. Indeed, it has been found that the diffusion constant measured from the pulse width with the coupling field on was $240 \text{ cm}^2 \text{ s}^{-1}$, as opposed to the value of approximately $30 \text{ cm}^2 \text{ s}^{-1}$ mentioned above without the coupling field present during storage [280].

5.3.5 Four-Wave Mixing Effects

Though a quantum memory candidate in its own right, a detrimental effect that can occur for other quantum memories in warm vapour, as well as other systems, is four-wave mixing (FWM). This is very similar to the spontaneous Raman scattering mentioned above, however, in FWM it is *coherent* Raman scattering. This process can be described by a double- Λ system, such as the ones shown in Fig. 5.4 for (a) an EIT-like system and (b) a Λ -GEM-like system. This scheme leads to the creation of an off-resonance Stokes field \mathcal{E}_s . The most common form of FWM is coherent anti-Stokes Raman spectroscopy (CARS), or degenerate FWM, where two coupling field photons of frequency ω_c interact with one photon of the probe field with frequency ω_p [281]. The equations for momentum and energy conservation are then given by

$$\omega_s = 2\omega_c - \omega_p \quad (5.14)$$

$$\mathbf{k}_s = 2\mathbf{k}_c - \mathbf{k}_p, \quad (5.15)$$

where ω_s is the frequency of the Stokes field and \mathbf{k}_s is its wave-vector. For the EIT-type system shown in Fig. 5.4(a) this means that $\omega_s = \omega_p - \Delta_{hfs}$.

The generated, or perhaps seeded, Stokes field will interact with the probe field via the atoms as they propagate through the medium. The equations of motion that govern this co-propagation are similar to those for the probe propagation described in Sec. 2.5.3 and, for EIT, are presented in Ref. [282]. Through both theory [283, 279] and experiment




Figure 5.4: Four-Wave Mixing Double- Λ Scheme. Λ -systems created on the D1 line of ^{87}Rb with probe (P - solid red line) and coupling (C - thick red line) field (a) on resonance and (b) off resonance, as well as secondary Λ -system formed with the creation of a Stokes (S - dashed pink line) field. Δ_{hfs} is the ground state hyperfine splitting, Δ is the one-photon detuning of the probe and coupling fields, δ is the two-photon detuning and $\Delta_s = \Delta - \Delta_{hfs}$ is the one-photon detuning for the coupling-Stokes system.

[284, 285, 124, 278] it has been shown that at high OD, FWM in EIT is significantly enhanced, leading to effects such as amplification due to interaction with the Stokes field (and the corresponding added noise that accompanies this [195]), as well as loss of probe field as it is converted into the Stokes field [285].

In Λ -GEM, FWM is not as much of an issue due to the already-detuned first Λ -transition. This large one-photon detuning, and the broadening of the absorption line for the gradient reduce the optical depth, which FWM is sensitive to. Also the fact that GEM is an absorptive memory, as opposed to the transmissive nature of EIT, means that the probe field will be absorbed, again reducing FWM effects. Due to this, we would mainly expect to see FWM effects at the edges of the broadened Raman line, away from the probe centre frequency. FWM effects in Λ -GEM have been studied and are presented in more depth in Ref. [252]. One important point uncovered from this work is that circular polarisations for the probe and coupling fields are better for suppressing FWM than linear polarisations. This is because, without a constant magnetic field in the correct direction, linear polarisation will be a combination of left- and right-circular light, giving the light opportunities to interact with more excited states.

5.3.6 Magnetic Field Effects

Finally, we will discuss the issue of magnetic-field-induced dephasing. The issue of having a non-zero magnetic field is a prime concern in EIT experiments. For instance, nearly all the warm vapour experiments presented in Sec. 5.2 used some form of magnetic field shielding, normally in the form of μ -metal, to protect against constant magnetic fields. Both N. B. Phillips et al. [124] and K. F. Reim [7] reported that their memory lifetimes

Rate	Value	Comments
$\Gamma_{D,\Delta}$	$2\pi \cdot 160$ MHz	$\ll \Delta$
$\Gamma_{D,\delta}$	$2\pi \cdot 8$ kHz	
$\Gamma_{col,\Delta}$	$2\pi \cdot 8.5$ MHz	$\ll \Delta$
Γ_{diff}	$2\pi \cdot 3$ kHz	
$\Gamma_{P,\Delta}$	$2\pi \cdot 15$ kHz	$\ll \Delta$
$\Gamma_{P,\delta}$	$2\pi \cdot 3.5$ kHz	
$\Delta\omega_{diff}$	$2\pi \cdot 8$ kHz	\ll pulse bandwidth, for $t_s = 100 \mu s$
$\Delta\omega_\eta$	$2\pi \cdot 10$ kHz	\ll pulse bandwidth, for multi-element coils

Table 5.1: Decoherence Rates and Frequency Changes in Warm Vapours - approximate values for the warm vapour Λ -GEM experiment.

were limited by magnetic-field-induced decoherence. This is also an issue in cold atom experiments, where three pairs of magnetic field coils are used to attempt to compensate the constant (and sometimes gradient) magnetic fields present.

If all internal m_F states of a given F level are to be used for storage then, if there is any non-zero magnetic fields present, they will all be at slightly different energies. This will translate itself to a broadening of an EIT or Raman line. During storage, with slightly different energies, the Bloch vectors will precess at slightly different speeds around the Bloch sphere, leading eventually to a loss of coherence between the m_F states and therefore memory efficiency. A thorough analysis of this process is presented in the supplementary information of Ref. [142].

In Λ -GEM the situation is very different. We require a constant offset magnetic field, we also use only one m_F state of each F level, and magnetic field gradients are required to rephase the system. However, an asymmetric magnetic field reversal will lead to a situation similar to longitudinal diffusion, with frequency blurring and inefficient recall. To determine the order of magnitude of this effect we can compare it to the deviation of a desired field with the actual field. This will be shown in Ch. 7 for a multi-element coil, where precision control over the gradient to tens of kilohertz was achieved. We can therefore estimate the change of frequency due to non-symmetric magnetic fields to be

$$\Delta\omega_\eta \approx 2\pi \cdot 10 \text{ kHz.} \quad (5.16)$$

5.4 Conclusions

In this chapter we have presented a comprehensive literature review of experiments that have used warm vapour cells to perform light storage experiments. These included the work already undertaken with warm-atom Λ -GEM, showing efficiencies up to 87% and up to 20-pulse storage. We have also presented a range of experimental issues encountered with warm atomic systems, some exclusively to non-zero-velocity ensembles and some not. A summary of the size of these various effects is presented in Tab. 5.1. As can be seen, the major decoherence effects are transverse diffusion and coupling-field-induced power broadening, limiting coherence times to approximately $50 \mu s$, in good agreement with measurements. The effect of the change in frequency of a given atom due to either longitudinal diffusion or non-symmetric magnetic field switching is not entirely clear. However,

the fact that the change is small compared to the pulse bandwidth, as well as the length of the spinwave indicate it will be minor.

Digitising GEM

Oscar: “Gentlemen, we can rebuild him. We have the technology.”

- *The Six Million Dollar Man* (opening titles), ABC (America), 1974-1978, TV Series.

6.1 Introduction

In Ch. 4 we investigated the use of digital locking and control on a Schrödinger cat experiment. In this chapter we will apply the same concepts to the Λ -GEM experiment. We first present the Λ -GEM experimental set-up in Sec. 6.2. We then show how digitising the detection for the Λ -GEM experiment is useful for data acquisition in Sec. 6.3, used here to collect noise measurement data. Following this, we detail how control of the system was changed to field programmable gate arrays (FPGAs) from a less flexible digital system, and discuss the benefits this will have for further experiments in Sec. 6.4. Finally, in Sec. 6.5 we present another use of digital locking for the Λ -GEM experiment - to perform digital demodulation of the heterodyne signal.

Some of the work presented in this chapter has been published in the journal article:

- M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler, *A Room Temperature Quantum Memory*, *Nature Physics* **7**, 794 (2011).

6.2 Λ -GEM Experimental Set-Up

The set-up of the warm rubidium-87 vapour Λ -GEM experiment is shown in Fig. 6.1. We used a Coherent MBR titanium:sapphire (Ti:Saph) laser with 1.2 W at 795 nm, pumped with a 532 nm Coherent Verdi V10 running at 8 W. We used a half wave-plate and a polarising beam-splitter (PBS) to send some of the laser power through an EO-Space fibre-coupled electro-optic modulator (EOM), driven at 6.8 GHz. This drove the light out of the carrier and into sidebands at $\pm n \cdot 6.8$ GHz sidebands (see inset (ii) of Fig. 6.1). Another EOM (NewPort free-space model 4002M) was used to create a Pound-Drever-Hall (PDH) locking modulation (as discussed in App. C). A ring cavity with a finesse of 100 (Cav1) was placed after the EOM. This was used to lock to the +6.8 GHz sideband, therefore filtering out all other sidebands. The laser was blue-detuned approximately 1.5 GHz from the D1 $F = 2 \rightarrow F' = 2$ transition, determined by viewing fluorescence through a reference Rb vapour cell using a video camera. The laser frequency was locked here using the Ti:Saph’s internal reference cavity.

It has been found previously that optimal efficiency occurs when the probe and coupling fields have the same circular polarisations [4]. Therefore polarisation optics, such as PBSs,




Figure 6.2: Zeeman States, used for Λ -GEM to create quasi-two-level atom on the D1 transition with: P - probe field; C - coupling field; Δ - one-photon detuning; dashed lines represent Zeeman splitting of ground states. Here we assume memory operation on the $m_F = 1$ states, with σ^+ polarisations for the probe and coupling fields. Energy levels not to scale.

could not be used for combining and splitting the two fields. Instead, we used a second ring cavity (Cav2). The LO field was split on a third PBS to create a locking field. This beam passed through an EOM (to again create a PDH locking modulation) and an acousto-optic modulator (AOM) from Opto-Electronic. The AOM was used to shift the frequency of the locking field to ensure that Cav2 was doubly resonant with the probe field in the TEM_{00} mode and the locking field in a TEM_{01} mode (see inset (ii) of Fig. 6.1). The reason for the difference in modes was to remove any leakage of the locking field from the heterodyne detection. The reason that the probe was not used to lock Cav2 was that it needed to be pulsed for memory operation. The coupling field's mode and frequency meant that it would be reflected off Cav2, with approximately 500 mW of power.

The co-propagating probe and coupling fields were then passed through a half and quarter wave-plate to set their polarisations going into the memory, consisting of a 20-cm-long cell of isotopically pure ^{87}Rb and 0.5 Torr krypton buffer gas with a diameter of 25 mm, heated to approximately 80°C with a filament heater. The cell was surrounded by a pair of 40-cm-long variable-pitch solenoid coils to create the input and output magnetic field (i.e., Zeeman) gradients for GEM. These, in turn, were surrounded by another solenoid of constant pitch to create a dc offset magnetic field and, around this, were a couple of layers of μ -metal to shield the memory from external constant magnetic fields. The magnetic field gradients created were on the order of 25 kHz cm^{-1} and, by switching off one coil, this could be reversed in approximately 500 ns. The dc offset was approximately 6 G.

This offset magnetic field led to Zeeman-splitting of the m_F states of the two F levels (see App. A). As we require a magnetic-field-induced frequency gradient for memory

operation we cannot use the $m_F = 0$ states. Therefore we used either the $m_F = 1$ or $m_F = -1$ states of both F levels by correctly selecting the polarisations and frequencies of the probe and coupling fields. This created our quasi-two-level Λ -GEM system, allowing us to use the long-lived coherence between the $F = 1$ and $F = 2$ levels of ^{87}Rb for the memory (see Sec. 3.4.4). This is illustrated in Fig. 6.2. With the ground-state Zeeman-splitting of $\pm 0.7 \text{ MHz G}^{-1}$ for the $F = 1, 2$ ground levels of ^{87}Rb (see App. A) we obtain a frequency-splitting of 1.4 MHz G^{-1} . To ensure as much of the atomic population as possible resides in the $F = 1$ ground level, the coupling field was turned on approximately $10 \mu\text{s}$ before the probe pulse was sent into the memory. If this was not the case, then the memory efficiency would decrease and absorption of the coupling field inside the memory cell would occur.

On leaving the memory, the probe and coupling fields passed through a 7.5-cm-long cell with a natural mixture of ^{87}Rb and ^{85}Rb , heated to 140°C . Due to the frequencies used for the probe and coupling fields, the coupling field will be absorbed by the ^{85}Rb while the probe will not (see Fig. A.2). This led to 60 dB of coupling field suppression, and 70% of the probe light being transmitted through the cell.

To be able to perform phase- and mode-sensitive detection we used a heterodyne detection scheme, as discussed in Sec. 2.3.5. We created the local oscillator (LO) field on a PBS after Cav1. After the filtering cell the probe and LO were combined for the heterodyne detection, with a visibility of roughly 97%. The detectors used in this experiment were silicon photo-diodes (model S3883 from Hamamatsu), having an efficiency of approximately 95% at 795 nm and a bandwidth of between 10-15 MHz. The relative phase between the probe and LO could not be locked as the probe was pulsed and therefore there could be no constant error signal.

The coupling, probe and local oscillator beam paths all contained single-pass AOMs to allow for fine control over the relative frequencies (here between 80-90 MHz) and temporal gating triggered by MiniCircuit switches from a National Instruments (NI) BNC-2110 digital and timing input/output board. The probe AOM driving signal could also be mixed with an envelope function from an Agilent arbitrary-waveform generator to create Gaussian pulses before the switch. This was again triggered by the NI board, as were the currents to the variable-pitch solenoids and the memory heater via solid state relays from VSHolding. The locking of the two ring cavities was initially provided by the analog proportional-integral (PI) controller characterised in Sec. 4.4.3.

6.3 Digital Acquisition for Noise Measurements

In previous chapters we have discussed the suitability of GEM, and in particular Λ -GEM, as a quantum memory candidate due to the high efficiencies achieved. However, efficiency alone cannot determine if a memory can operate in the quantum regime. If the memory were to somehow amplify the stored information, for instance, a high efficiency would be measured but the output pulse would contain excess noise. With Λ -GEM we must be especially vigilant as the strong coupling field has the potential to add noise to the output state due to spontaneous Raman scattering (discussed in Sec. 5.3.4).

One way of testing whether a memory is operating in the quantum regime is to store a fundamentally quantum signal such as squeezing [238, 239] or entanglement [26] and measure the degradation of the state. In the absence of these signals, noise measurements can be performed using quantum process tomography [287, 288, 289, 290, 291], which

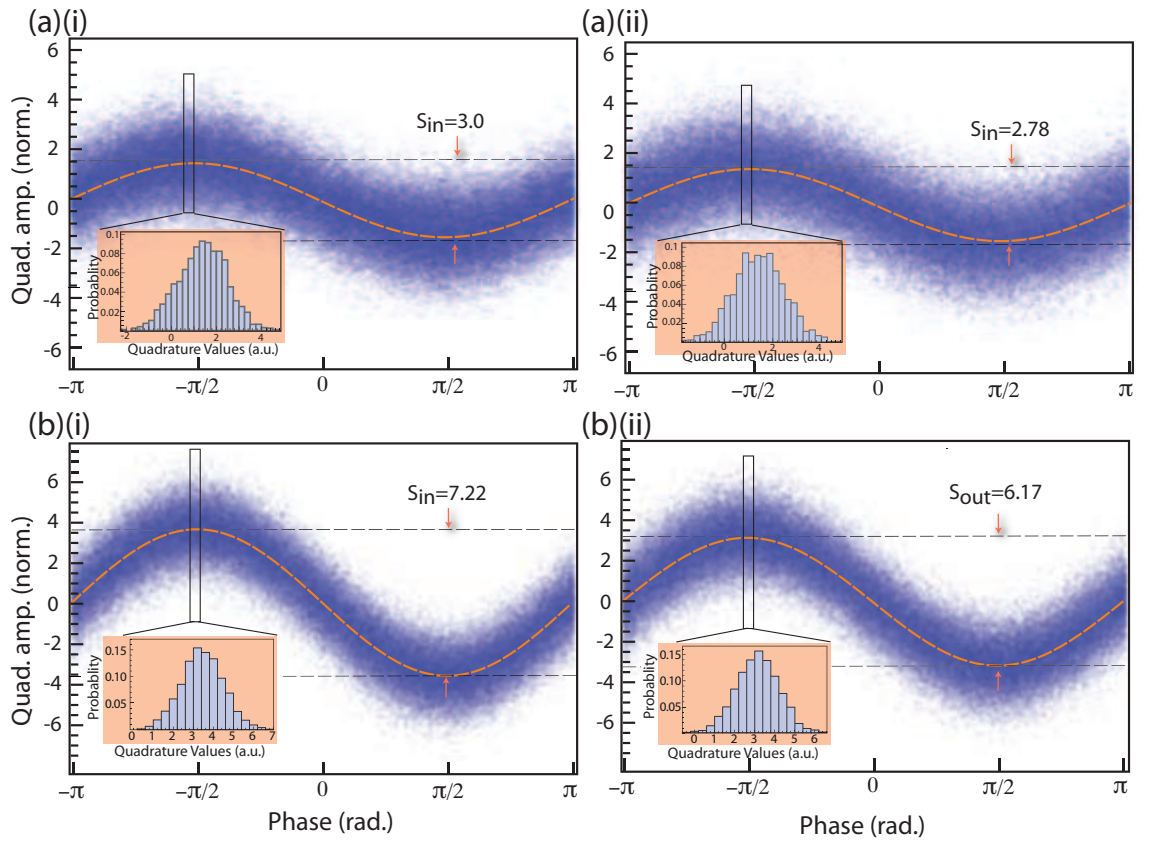


Figure 6.3: Λ -GEM Noise Measurements 1. (a) Normalised quadrature amplitudes as a function of relative phase between probe and LO fields for (i) input with mean photon number $n_p = 0.67$ photons and (ii) corresponding echo quadrature amplitudes. S_{in} is the input amplitude and S_{out} is the output amplitude. (b) The same for $n_p = 3.4$. Insets show histograms of the quadrature values for the phase region denoted by the black rectangle. Adapted from Ref. [286].




Figure 6.4: Λ -GEM Noise Measurements 2. (a) Fidelity of pulses as a function of their mean photon number n_p . The dashed line denotes the no-cloning limit. (b) T-V diagram, with shapes corresponding to the same n_p as in (a). Shaded region corresponds to no-cloning regime. Adapted from Ref. [286].

can characterise the response of a memory for any input state. Even without process tomography, coherent states can be used to compare the memory to important quantum benchmarks. As discussed in Sec. 3.2.1, if we could show storage with a fidelity over 50%, we would be performing better than any classical memory [96], and if we can do the same over 68% then we would be in the no-cloning regime [97]. Weak coherent states with mean photon numbers $n_p < 1$ through to 10^4 were used to determine that two-level GEM in a solid state system did not add noise and, with a storage efficiency of 69%, therefore broke the no-cloning limit for $n_p < 30$ [181].

We performed a similar experiment with the warm ^{87}Rb vapour Λ -GEM system. To perform noise measurements it is necessary to scan the phase of the LO relative to the probe. Here, the fact that we couldn't lock the phase between the two allowed us to scan the phase “for free”. However, we needed a phase reference to be able to accurately associate a phase with the input and echo pulses. To do this we sent in a strong pulse of probe light, 3 μs in length, without the coupling field on. The weak 2- μs -long input pulse was sent in 9 μs after the reference pulse, close enough that there would be no phase difference between the two. The input was stored for 3 μs , with the coupling field turned off for 1 μs during storage to reduce the coupling-field-induced scattering (see Sec. 5.3.4).

To acquire statistically significant data for all quadratures with no control over the relative phase meant that a lot of data needed to be taken. As with the data collection in Sec 4.5, digital acquisition with a NI PXI-5124 data acquisition card was used to allow for fast collection. In this case a hardware trigger from the NI BNC-2110 device was used to synchronise the data collection to the start of the run (which lasted less than 20 μs in total), with many runs being saved to the one file to minimise data transfer time. In total 100,000 input traces and 100,000 output traces were saved for five values of n_p varying from $n_p = 0.67 \rightarrow 22.4$.

The integrated amplitude of the inputs and the echoes were then processed and plotted as a function of the relative phase. Figure 6.3 shows these traces for two values of n_p . Firstly it can be seen that the memory efficiency was approximately 86% for $n_p = .67$ and 73% for $n_p = 3.4$ (remembering efficiency is proportional to intensity and not amplitude). Also, from the histograms it appears that very little noise was being added by the memory.

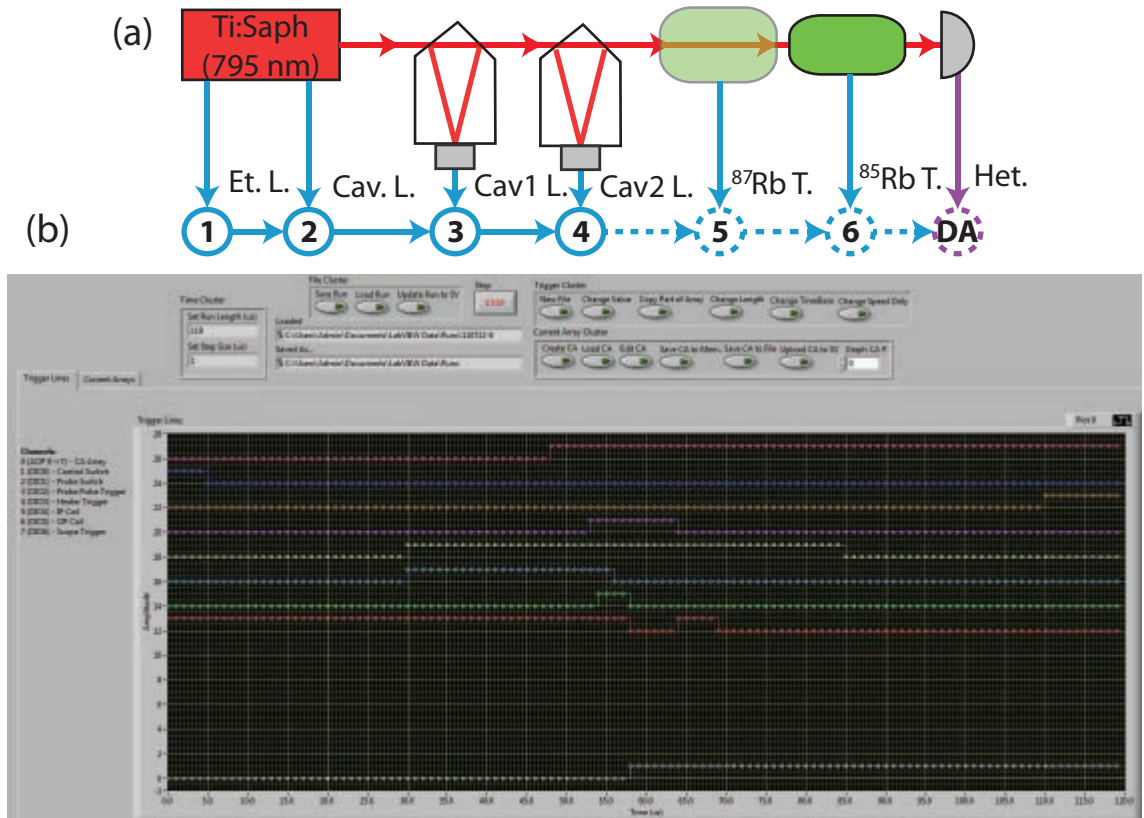


Figure 6.5: FPGA Control and Monitoring. (a) Flow diagram showing signals monitored by the FPGA-based locking system, with: Ti:Saph - titanium:sapphire laser; Et. L. - etalon lock; Cav. L. - internal reference cavity lock; Cav1 L.- first cavity lock; Cav2 L. - second cavity lock; $^{87}\text{Rb T.}$ - memory cell temperature; $^{85}\text{Rb T.}$ - filter cell temperature; Het. - heterodyne signal (sent to digital acquisition - DA - device). Also shown is the ordering for sequential locking, arrows and numbers indicate the direction, dashed lines indicate components that have not yet been implemented. (b) Screenshot of the FPGA-based triggering system used for coupling field on/off; probe on/off; probe pulse triggering; heater on/off; input and output coils on/off; and triggering the oscilloscope used to monitor observe the memory operation. This triggering system also controls analog output timings for Ch. 7, and has a “batch” running mode, as described in Ch. 8.

To quantify the performance of the memory, these quadrature values were used to reconstruct the density matrices for the various n_p coherent states, using an iterative algorithm for maximum-likelihood estimation [292]. From the density matrices the fidelities of the states could be determined, as well as a T-V diagram (see Sec. 3.2.1). Both characterisation methods are shown in Fig. 6.4. As can be seen, for $n_p \leq 16.3$ the memory performed above the no-cloning limit, with a maximum unconditional fidelity for the $n_p = 0.67$ state of 98%. However, it was found that there was some noise added to the output, due mainly to instabilities in the analog locking system. For more information on these measurements and the results, see Ref.s [258] and [286].

6.4 FPGA Control System

The set-up used in the previous section, where a NI BNC-2110 device was used to create hardware triggers for all components and analog PIs were used to lock the two cavities, was not optimal for digital acquisition. For instance, digital control over the two cavities could have been implemented for greater lock performance, as per Sec. 4.4.2, or the code could have been altered to alert the data acquisition program when the cavities were becoming unstable and these points could have been removed from the data.

We therefore decided to change the system over to complete FPGA control. The first step was to swap the locking of the cavities over to the digital locking code developed in Ch. 4. The main difference here was that we used signals which were demodulated outside the FPGA, due to a lack of high-speed analog-to-digital converters. Sequential locking logic was used between Cav1 and Cav2, however total automatic locking could not be used for Cav1 due to the nature of the input: namely that there were two peaks with the same height (both plus and minus 6.8 GHz sidebands, see inset (ii) Fig. 6.1). Therefore the peak search function described in Sec. 4.3.2 could not distinguish between the two. Extra *Do Not Relock* logic was then added to the code that allowed for initial automatic locking but, once the system became unlocked, would not relock until the user enabled the automatic locking function again. The idea behind this is that the user would reduce the scan range so that only the sideband of interest can be viewed, then instigate locking followed by activation of the *Do Not Relock* function. If the system came unlocked, the user would again ensure that only the correct peak is present in the scan range, remove the *Do Not Relock* logic, and then reactivate it once the system is locked. The ability to have automatic sideband-selection will be available in the near future.

The spare analog inputs on this FPGA were not wasted but put to work monitoring other parts of the system. For instance, the etalon and internal reference cavity locking signals from the Ti:Saph were monitored here. Though the FPGA could not control the laser locks, it could be used to monitor them, and basic threshold logic was used to determine if the laser was locked or not. This was a handy feature as it was not always apparent that it was the Ti:Saph that had come unlocked if the memory stopped working. These laser thresholds could also be incorporated into the sequential locking logic. The other parameters that were monitored were the temperatures of the memory and filter cells. Temperature probes provided a voltage level proportional to how hot the cells were, and a rough calibration was used inside the LabVIEW Real Time (RT) code to display approximate temperatures on the screen. This was again handy for monitoring all system parameters in one place, and if one were so inclined, code could be written so that temperature thresholds could be set to ensure that data would only be acquired when the cells were in the correct temperature range (or potentially even feedback control applied). This is shown in Fig. 6.5(a).

To allow for greater flexibility in triggering the system, an FPGA control code was also written for creating and implementing memory runs. Unlike the previous control, which was a user-defined array of zeros and ones and a timing base, the new FPGA-based control code allowed users to change individual values one-at-a-time (as before) but also to change a range of values at once, copy part of one output to another, as well as change the time-base of the entire run. It also displayed the runs on screen and allowed for the runs to be saved and reused later. To allow for the greatest amount of manual optimisation, all frequencies and magnetic fields were controlled externally to this code. On top of this, the triggering array could be altered while running to optimise timings through the RT

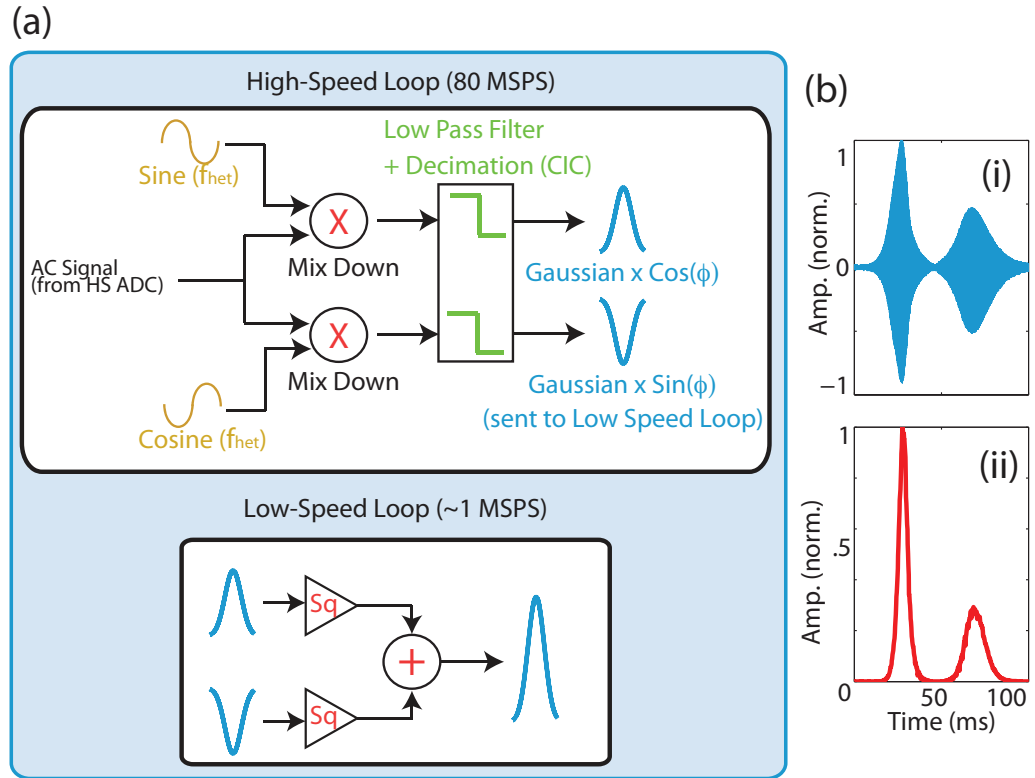


Figure 6.6: Digital Phase-Insensitive Heterodyne Demodulation. (a) Schematic of the FPGA code where both a sine and cosine at the heterodyne beating frequency f_{het} are used to demodulate the input, followed by a low pass cascaded-integrator-comb (CIC) filter in the high speed loop. These two Gaussians are then sent to the low speed loop where they are squared (Sq) and then added to give the Gaussian intensity profile. (b) Results showing (i) input modulated at 6 MHz and (ii) demodulated, squared output. Both pulses are normalised to maximum signal size.

code. A screen-shot of this code is shown in Fig. 6.5(b).

The extra flexibility provided by this FPGA code proved useful for the work carried out in Ch.s 7 and 8.

6.5 Digital Phase-Insensitive Heterodyne Demodulation

In this section we present one final use for digital data acquisition for the Λ -GEM experiment. As mentioned in Sec. 6.2, the relative phase of the probe and LO could not be locked. This meant that we had to use heterodyne detection as opposed to homodyne detection. This was because heterodyne detection continuously shows the full pulse envelope, though modulated, while homodyne detection does not. For optimisation purposes it is always better to be able to view the full envelope and, if possible, the intensity profile rather than the amplitude to give a better impression of the efficiency by eye. However, we do not want to sacrifice the added information, selectivity, or sensitivity of heterodyne detection.

To overcome this issue we implemented digital, phase-insensitive demodulation of the

heterodyne signal. A schematic diagram of the code is shown in Fig. 6.6. The idea here is very similar to PDH locking (see App. C) in that the input signal is sent to an FPGA via a high-speed analog-to-digital converter and demodulated by firstly multiplying it with a sine wave at the same frequency and then low-pass filtering it. In our case the frequency was the frequency difference between the probe and local oscillator f_{het} and, unlike PDH locking, the phase of our demodulating sine was not fixed relative to the modulation. However, if we perform the same operation with a cosine as well, and then square both signals separately in the second, low-speed, loop, we end up with $Gs^2[t] \cdot \text{Cos}^2[\phi(t)]$ for the sine-modulated signal and $Gs^2[t] \cdot \text{Sin}^2[\phi(t)]$ for the cosine-modulated signal, where $Gs[t]$ is the Gaussian amplitude envelope function and $\phi(t)$ is the phase between the heterodyne input and the sine signal.

As long as $\phi(t)$ does not vary much within one demodulation cycle (i.e. $\phi(t) \approx \phi$) then, if we add both terms together, we will obtain $Gs^2[t]$ as $\text{Sin}^2[\phi] + \text{Cos}^2[\phi] = 1$. With this information we could determine the efficiency in real-time using the LabVIEW code to help with the optimisation process. An example of the phase-insensitive demodulation is shown in Fig. 6.6(b), where a pulse modulated at 6 MHz is sent into the FPGA and digitally demodulated to produce the intensity envelope. This code could be extended to provide real-time feedback on the memory efficiency.

6.6 Conclusions

In this chapter we implemented digital control and locking schemes on the Λ -GEM system to improve the experimental set-up. With the implementation of digital acquisition we were able to take one million traces for noise measurements to show that Λ -GEM breaks the no-cloning limit for less than 22 photons per pulse, with a maximum unconditional fidelity of 98% for a pulse with a mean photon number of 0.67. We also detailed the digital locking code that will be used to control the system for the experiments discussed in Ch.s 7 and 8, and how digital programming can be used to digitally demodulate a heterodyne signal, with potential extensions to real-time efficiency monitoring.

Precision Spectral Manipulation Using GEM

Carter: “I know that you have a device that can...create a wormhole, or, uh, bend time, or make you invisible... or a wormholing, time-bending, invisibling device, that shields you from the mind.”

Stark: “Yes, he said ‘invisibling’.”

- “Before I Forget” *Eureka*, SyFy, 1 August 2006, TV Series.

7.1 Introduction

In previous chapters we have discussed the importance of quantum memories, as storage devices, for quantum information processing, and investigated the ability of GEM to reproduce, as reliably as possible, the stored information in question. However, in addition to memories, a method for conditioning qubits in terms of their timing, frequency and bandwidth would be another useful component of a quantum communication system. The ability to coherently manipulate the spectrum of pulses would prove a key tool for allowing quantum information transfer between systems with different bandwidths. This ability could lead to increased bit rates over quantum communication channels [293]. Being able to alter a pulse’s shape, as well as its bandwidth, could also lead to increased bit rates due to a decrease in losses caused by pulse aberrations through various media (i.e., optical fiber [294]). While being able to controllably alter a pulse’s phase could allow for the creation of gates [295, 296], or states [69, 70, 71], for quantum computing.

There are numerous options available for conditioning operations, some of which would be suitable for quantum information applications. Diffraction techniques [297], electro-optic modulators (EOMs) [298, 299, 300], and acousto-optic modulators (AOMs) [301] have been used to modulate pulse amplitude and phase, as well as pulse compression and shaping. Spatial-light modulators (SLMs) can also be used to shape both classical [302] and quantum [303, 304, 305] pulses. Another classical tool is temporal lensing systems [306], which can be used to compress [307, 308] and expand [309] pulses, as well as perform time-to-frequency conversion [310, 311]. Four-wave mixing can also be used to alter the frequency of a pulse (for instance, see Ref.s [312, 313]). As well as these, there are quantum techniques that can be used to manipulate pulses. For instance, three-wave mixing [314], quantum pulse gates [315, 316], and pulsed frequency up- [317, 318, 319], and down- [320, 321], conversion have all been shown to be able to coherently alter either the temporal or spectral profile of optical pulses. Frequency up-conversion has also been used as the basis of an optical switch [322].

Given that a memory will already be a part of a quantum information network, as information will need to be stored, an enticing possibility is that qubit conditioning could be built into the memory functionality to allow for purification and distillation operations [323, 324], among other tasks. If this is possible, it would not only save the addition of extra components (and the extra loss and complexity this entails) but would also allow dynamic programming of the manipulation and timing - as the information would not simply be passing through the system, but be trapped. It could also allow for controllable multiplexing within the memories - a powerful tool where multiple signals are bundled into one over a communication channel, increasing bit rates by orders of magnitude. Multiplexing could be achieved with, for instance, different spatial [104], temporal [105] or frequency modes in a quantum memory. Being able to perform multiplexing with a quantum memory has also been suggested as a way of improving quantum repeater designs by speeding up the entanglement generation process by orders of magnitude [105], and reducing their sensitivity to memory coherence times [106].

Electromagnetically-induced transparency (EIT)-based memory has been investigated with respect to pulse shaping and splitting inside a coherent memory [257, 325]. Angular and frequency multiplexed storage of information has also been demonstrated using EIT with cold atomic ensembles [326, 327, 328]. By introducing a second, counter-propagating, coupling field, an electromagnetically-induced grating can be created, which allows for optical switching and routing [329], as well as positioning [330]. Manipulation of stored information has also been demonstrated in photon echo-based quantum memories. For instance, the ability to perform spectral manipulation using a modified photon echo scheme has been theoretically shown [331] and, using the three-pulse photon echo, interference between two pulses has been experimentally demonstrated [332]. Four-wave mixing (FWM) has also been used to multiplex up to 12 spatially divided signals in a cold atomic ensemble [333].

Previous experimental work on Λ -GEM has shown that it is capable of storing up to 20 pulses simultaneously [4], and is able to manipulate stored light in a number of ways: it has been used to resequence pulses; stretch or compress the bandwidth of stored pulses [9]; add a frequency offset to the recalled light [4]; and interfere two pulses within the memory [10]. Modelling has shown that GEM schemes are capable of much more. For example, it could be used as an optical router [187] or all-optical multiplexer [188].

In this chapter we investigate, in detail, proposals that make particular use of the frequency-encoding nature of GEM to coherently manipulate the spectrum of stored pulses, filter modulated pulses, and combine or interfere pulses of different frequencies, as theoretically described in Ref. [185]. Using our memory we also demonstrate the feasibility of frequency multiplexing, where we can independently store and recall pulses of different frequency. To be able to perform these operations we need fine control over the input and output frequency gradients. This is achieved with the replacement of the variable-pitched solenoids, described the previous chapter, with a multi-element coil.

The remainder of this chapter is structured as follows: the experimental details of the multi-element coil are presented in Sec. 7.2. We then, in Sec. 7.3, present experimental results characterising basic frequency manipulation operations, as well as demonstrating frequency-domain engineering with precision control of magnetic field gradients provided by the multi-element coil. Finally, in Sec. 7.4 we will investigate the ability of GEM to controllably alter the phase of stored information using cross-phase modulation.

Most of the work presented in this chapter has been published in the journal articles:

- B. M. Sparkes, C. Cairns, M. Hosseini, D. Higginbottom, G. T. Campbell, P. K. Lam, and B. C. Buchler, *Precision spectral manipulation: a demonstration using a coherent optical memory*, *Physical Review X* **2**, 021011 (2012)
- M. Hosseini, S. Rebić, B. M. Sparkes, J. Twamley, B. C. Buchler, and P. K. Lam, *Memory-enhanced noiseless cross-phase modulation*, *Light:Science & Applications* **1**, e40 (2012).

7.2 The Multi-Element Coil

The experiments presented in the next section used the same set-up as that presented in Sec. 6.2 and the digital control system from Sec. 6.4, with one main difference: to achieve fine control over the magnetic field gradient, required to perform the manipulation operations, it was necessary to move away from the two variable-pitch coils. Instead, a set of eight solenoids, with four turns each, were placed approximately 1 cm apart and spanning 30 cm along the memory (of length 20 cm). This is shown in Fig. 7.1(a). By sending different currents to the individual solenoids of this multi-element coil (MEC) and using the superposition principle for magnetic fields, i.e., $B_{tot}(z) = \sum_j B_j(z)$, a variety of different gradients could be produced. Due to the larger diameter of the new coils, we were not able to use the dc magnetic field offset solenoid.

To determine the correct currents for a desired magnetic field, a Mathematica code was created (see App. D), with the currents determined from the fitting being exported as a text file. Once the currents for a given gradient were determined, the values were entered into a LabVIEW Real-Time (RT) code, which transferred the data to an National Instruments (NI) field programmable gate array (FPGA) 7852R device. The eight analog output voltages, corresponding to the desired currents, were sent to a current driver (built in-house) that created a current proportional to the input voltage. Finally, these currents were sent to the respective individual solenoids to create the desired gradients.

These gradients will be discussed in the following sections, but included various linear gradients with various offsets, as well as different linear gradients for different parts of the memory (splitting in to either two or three parts). Some examples of these gradients are shown in Fig. 7.1(b), plotted in terms of the two-photon detuning δ (see Fig. 3.4). These gradients were measured by moving a Gauss-meter down the centre of the coils, and measuring the field in the z -direction (i.e., the direction of propagation) every 2 cm. This measure of the magnetic field was converted to a two-photon detuning via the following equation

$$\delta(z) = 2g_F \cdot B_{tot}(z) - \delta_o, \quad (7.1)$$

where $g_F = 0.7 \text{ MHz G}^{-1}$ is the Landé factor, and δ_o is an arbitrary two-photon offset, defined for a set dc magnetic field created with the MEC (650 mG) and probe and coupling field detunings. As can be seen from the figure, the desired, best possible, and measured magnetic fields agree very well for the first two examples but the magnetic field resolution breaks down for the third. When thinking of ways to improve the resolution, one option that springs to mind would be to include more coils in the MEC. However, after modelling the behaviour using the Mathematica code, it was found that adding more coils would not greatly improve the resolution, and would not be feasible due to other factors such as the maximum number of FPGA analog outputs (eight) and limitations on the maximum current that could be supplied to the current driver (up to 10 A each for the positive and

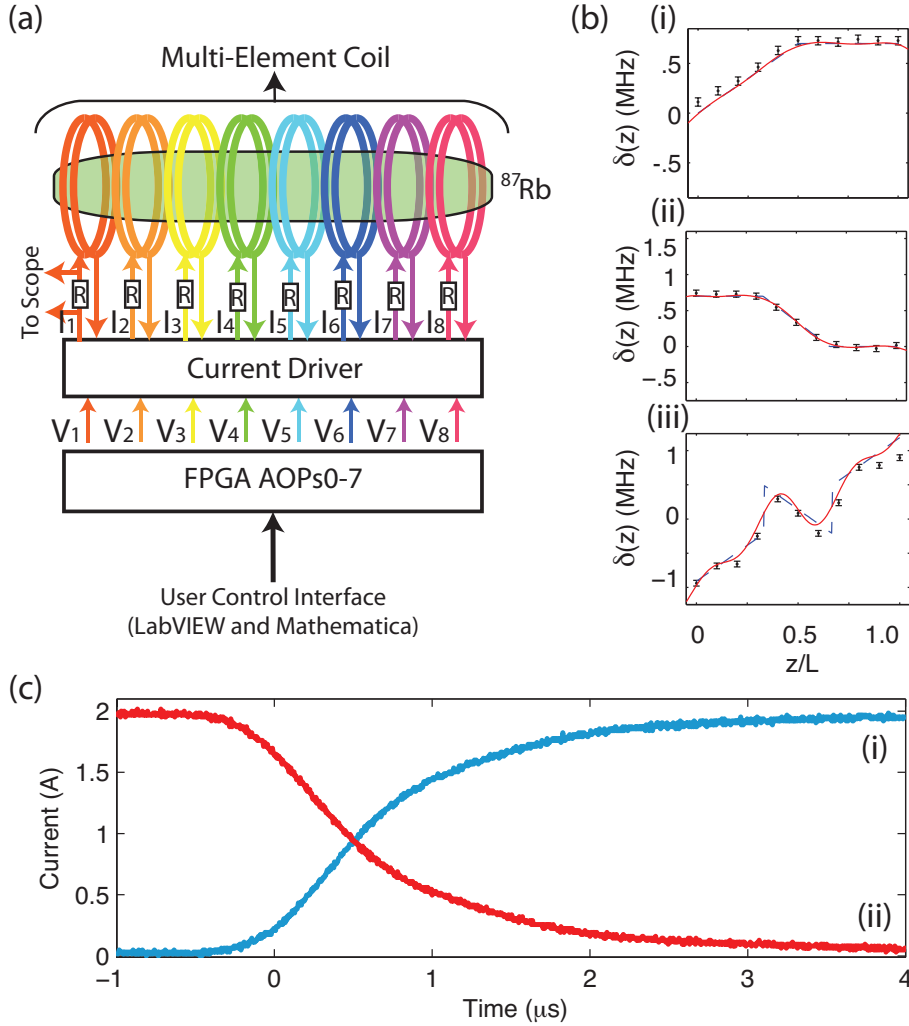


Figure 7.1: Multi-Element Coil Set-Up. (a) A schematic diagram showing the set-up of the eight solenoids comprising the multi-element coil, with: FPGA - field programmable gate array card; AOP - eight analog outputs with voltages $V_1 \rightarrow V_8$; $I_1 \rightarrow I_8$ - currents sent to coils; R - resistor placed in each circuit ($R=5\ \Omega$); and Scope - position of oscilloscope for measuring the currents in the coils. (b) Examples of two-photon detuning δ as a function of z -position along the memory (normalised to length L) for (i) two gradient components, (ii) three gradient components (continuous), and (iii) three gradient components (non-continuous). Blue (dashed) line corresponds to the desired field, red (solid) line corresponds to best possible fit given physical set-up, and points correspond to the measured magnetic field (error bars due to sensitivity of Gauss-meter). (c) Switching behaviour of coils with $R=5\ \Omega$ and 2 A both (i) on and (ii) off.

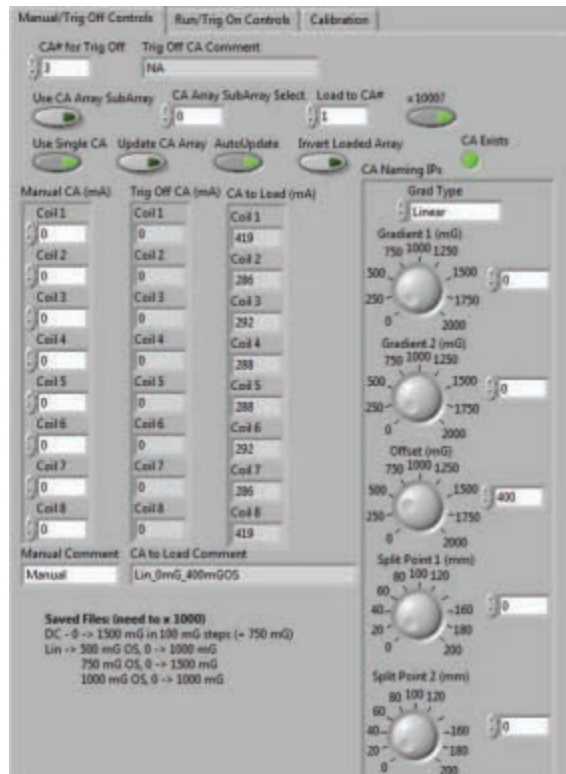


Figure 7.2: LABVIEW MEC Gradient Code. A screenshot of the LabVIEW code used to choose gradients when the memory was not being triggered. It could also be used to override gradients in the run code during triggering.

negative rails). With the eight-coil MEC it was not possible to create gradients with sharp discontinuities, and strategies were employed to minimise the use of these gradients.

Along with matching the desired and actual magnetic fields as well as possible, it was also important to minimise the switching time. Here we achieved this by placing a resistor in the system to damp oscillations caused by the inductance of the coils. It was found that the optimal resistance was approximately $R = 5 \Omega$ in series with the coils. The addition of resistors in the set-up also allowed the currents in the eight coils to be constantly monitored using two four-channel oscilloscopes, with the voltage drop across each resistor V measured and converted to a current I using the equation $I = V/R$. Figure 7.1(c) shows the switching time of the eight coils using this method. As can be seen, the minimum switching time was on the order of 2-3 μs , slower than the 0.5 μs switching time of the variable-pitch coils coupled with the previous triggering configuration. This was mainly limited by the triggering hardware's running speed of 750 kHz.

The FPGA control, presented in Sec. 6.4, allowed for synchronisation between the switching of the eight analog coil control voltages with the triggering of the other components required to operate GEM (i.e., probe pulse trigger, control switch, etc.). This would not have been possible with the old NI BNC-2110 control hardware. The FPGA code allowed for up to eight separate arbitrary gradients (each consisting of an array of eight currents, one for each MEC coil) for each run.

To allow for some of the manual flexibility of the variable-pitch solenoids, many different gradients were loaded from the Mathematica code to the LabVIEW RT code before the




Figure 7.3: Basic Λ -GEM Set-Up, with: P - probe field; C - coupling field; LO - local oscillator field; Cav - ring cavity; MEC - multi-element coil; ^{87}Rb - memory cell; ^{85}Rb - filter cell; BS - 50:50 non-polarising beam-splitter; D - detector.

experiment was undertaken. A screenshot of the LabVIEW RT code is shown in Fig. 7.2. The user could either manually enter values for each of the eight currents to create any desired field (left side), or they could pick a type of gradient (i.e., “Linear”), followed by the size and offset (right side). These gradients could be loaded into a run before undertaking an experiment. These could also be loaded over the top of the currents set in a run to allow for dynamic optimisation of the memory for each of the various operations detailed below. This code, and its ability to interact with the already-developed FPGA triggering code, is another demonstration of the power of digital control to improve quantum optics experiments.

Figure 7.3 shows the basic set-up for this experiment. As in Sec. 6.2, the probe and coupling fields, separated in frequency by 6.8 GHz (probe on the $F = 1 \rightarrow$ and coupling field on the $F = 2 \rightarrow \text{D1}$ transitions), were placed on Raman resonance 1.5 GHz blue-detuned from the D1 $F' = 2$ level. They were combined on a ring cavity before being sent into the 20-cm-long isotopically pure ^{87}Rb vapour cell, heater to 80°C. Upon retrieval, the echo and coupling fields passed through a filtering cell and the probe was detected using a local oscillator (LO) and heterodyne detection. For more details see Fig. 6.1.

7.3 Spectral Manipulation Experiments

In this section, we present the various spectral manipulation experiments undertaken with the MEC.

7.3.1 Centre Frequency Offset Characterisation

By adding an offset to the recall field, i.e., $\delta(z, t > \tau) = -\delta(z, t < \tau) + \delta_{os}$, the centre frequency of the echo ω_{co} relative to the input ω_{ci} can be altered and will be given by

$$\omega_{co} = \omega_{ci} \pm \delta_{os}. \quad (7.2)$$

The sign of δ_{os} is dependent on which m_F state is used for the memory. In this case it is the $|F = 1, m_F = 1\rangle$ state and therefore the sign will be negative (see App. A).

With the MEC we are able to vary the offset of the magnetic field down to tens of kilohertz precision, without altering the gradient. This fine control over the magnetic field

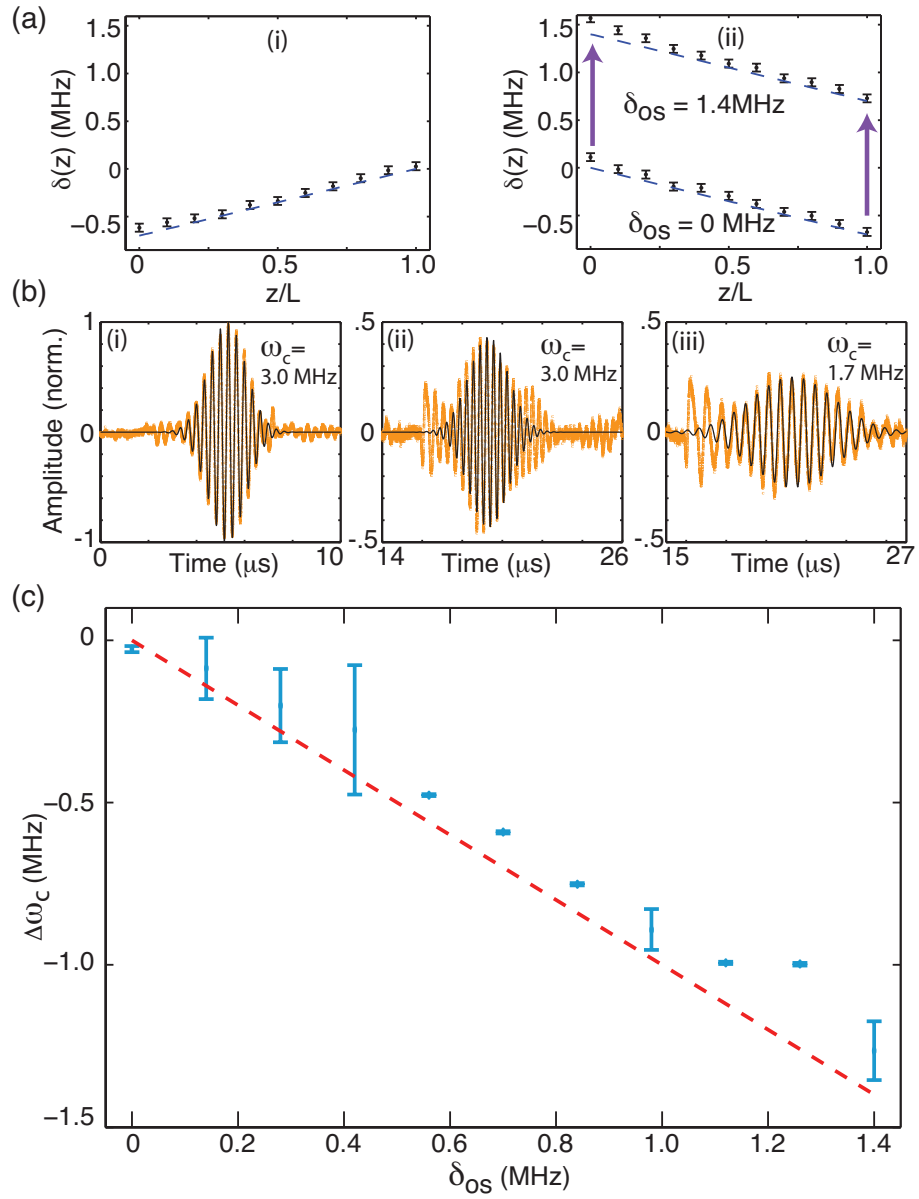


Figure 7.4: Centre Frequency Offset Characterisation. (a) Two-photon detuning δ as a function of position z along the memory (normalised to length L) due to (i) input gradient and (ii) output gradients with minimum and maximum gradient offset δ_{os} (as noted on figure). Blue (dashed) line corresponds to the desired field and points correspond to the measured magnetic field (error bars due to sensitivity of Gauss-meter). (b) Heterodyne data showing (i) input pulse; (ii) echo for recall with $\delta_{os} = 0$; and (iii) echo for recall with $\delta_{os} = 1400$ kHz offset. Orange points correspond to raw data, black lines correspond to modulated Gaussian fit to data, and ω_c values correspond to the centre frequencies of pulses extracted from the fits. (c) The change in centre frequency of the output pulses relative to the input pulse $\Delta\omega_c$ as a function of δ_{os} . Points represent the measured centre frequency (error bars from standard deviation of 100 traces), and the dashed line corresponds to the theoretical behaviour.

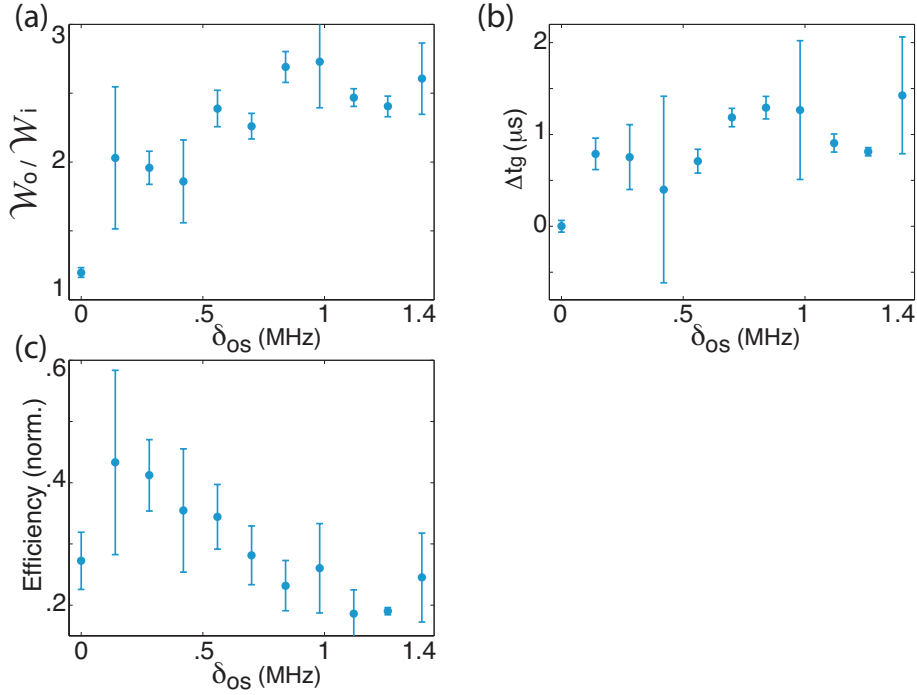


Figure 7.5: Extra Centre Frequency Offset Characterisation. (a) FWHM of the output pulses \mathcal{W}_o normalised to the FWHM of the input pulse \mathcal{W}_i , as a function of gradient offset δ_{os} . (b) Change in recall time Δt_g compared to recall with $\delta_{os} = 0$ as a function of δ_{os} . (c) Efficiency as a function of δ_{os} . In all figures points represent data and error bars are from standard deviation of 100 traces.

is shown in Fig. 7.4(a), where (i) the input and (ii) output gradients for varying magnetic field offsets are presented. As can be seen, the desired and measured magnetic fields are in close agreement.

Figure 7.4(b) shows single heterodyne traces for (i) input and (ii) echo with no applied offset, as well as (iii) echo with an applied offset of 1.4 MHz. A modulated Gaussian was then fit to the modulated data. This was of the form

$$\pm A \text{Cos}[2\pi\omega_c(t - t_o)] \cdot \text{Exp}\left[\frac{-(t - t_g)^2}{2\sigma_t^2}\right], \quad (7.3)$$

where A is the amplitude (the \pm was used to help fit the phase correctly in Mathematica), ω_c is the centre frequency (relative to the LO frequency), t_o is the shift from time $t = 0$ for the phase of the cosine, t_g is the centre position of the Gaussian in time, and σ_t is the Gaussian standard deviation¹. These fits were performed in Mathematica (see App. D) and, by performing them, we were able to obtain information about the amplitude, centre frequency, centre position etc. for each pulse. Examples of these fits are shown in Fig. 7.4(b) for the input, as well as the two echoes. Figure 7.4(c) shows a characterisation of the change in ω_c for a range of values of δ_{os} . This is compared with the behaviour expected from Eq. 7.2. As can be seen, the two are in good agreement.

While the frequency was shifted as we expect, there was also a substantial amount of

¹Related to the full-width-half-maximum (FWHM) of the Gaussian via $\mathcal{W} = 2\sigma_t\sqrt{2\ln[2]}$.

pulse stretching indicating that our pulses were subject to some dispersion in the memory. This was not surprising as the bandwidth of the memory was only slightly greater than the bandwidth of the pulse. Details of the dispersion in GEM can be found in Ref. [252] where the susceptibility has been explicitly calculated, clearly showing that near the edges of the absorption window there is great potential for slow light effects that will broaden recalled pulses. Having a larger memory input bandwidth would reduce the effect of dispersion on the pulse. The additional elongation for recall with a greater offset indicates that fringes of the magnetic field (i.e., those components that tail off at either end of the cell) may have affected the stored pulse, leading to greater dispersion. This is characterised in Fig. 7.5(a). As we will explain in Sec. 7.3.2, however, GEM provides a means to easily compensate for such pulse broadening.

It can also be seen from Fig. 7.4(b) that there was no significant change in storage time, as expected. Finally, the efficiency of the echoes as a function of gradient offset are shown in Fig. 7.5(c). As can be seen, as the pulse offset increased, and the edge effects became more pronounced, the efficiency decreased. This will be at least partly due to the fact that the gradient was no longer monotonic at the edges.

7.3.2 Input/Output Bandwidth Characterisation

By recalling with a steeper output gradient than input gradient, i.e., $|\eta(z, t > \tau)| > |\eta(z, t < \tau)|$ with τ being the switching time, the output bandwidth of the memory \mathcal{B}_o will be made greater than the input bandwidth. This change in bandwidth is, in turn, passed on to the echo as the absolute emission frequency relative to the centre frequency $|\omega(z) - \omega_c|$ of each atom along the ensemble will be greater, while the total excitation $\sigma_{12}(z, t)$ will remain unchanged. In this case, the output pulse will be compressed in time due to its now greater frequency spectrum. The opposite is also true, i.e., by recalling with a shallower gradient the output pulse bandwidth will be reduced and the pulse elongated in time.

The temporal profile of the output pulse, measured using the pulse full-width-half-maximum (FWHM) \mathcal{W}_o can be expressed as a function of the input profile \mathcal{W}_i and input/output gradient $\eta_{i/o}$ as

$$\mathcal{W}_o = \mathcal{W}_i \cdot |\eta_i/\eta_o|. \quad (7.4)$$

Again, it was the precision control of the MEC over the magnetic field that allowed us to alter the bandwidth by tens of kilohertz without altering the magnetic field offset. This is shown in Fig. 7.6(a), where experimental plots of (i) the input gradient, and (ii) output gradients with ratios $|\eta_o/\eta_i|$ from 1:1 to 3:1 are shown. Performing fits to individual pulses using Eq. 7.3 and the Mathematica notebook presented in App. D, as illustrated in Fig. 7.4(b), allowed for in-phase digital demodulation of the heterodyne data. This, in turn, allowed for averaging over many traces, something that would not have been possible with the non-demodulated data due to phase fluctuations between the probe and LO.

Figure 7.6(b) shows averaged demodulated input and output pulse amplitudes for different recall gradients. As predicted, the output pulses became more compressed as the recall gradient was increased. Though it did follow a linear relationship, it did not, however, follow Eq. 7.4, as can be seen from Fig. 7.6(c). This discrepancy is more evidence of the dispersion that is present when the pulse bandwidth is approximately equal to the memory bandwidth, as it was in our experiment. Numerical modelling of the experiment with Eq.s 3.34-3.35 and XMDS2 [334], using experimentally determined values for pulse

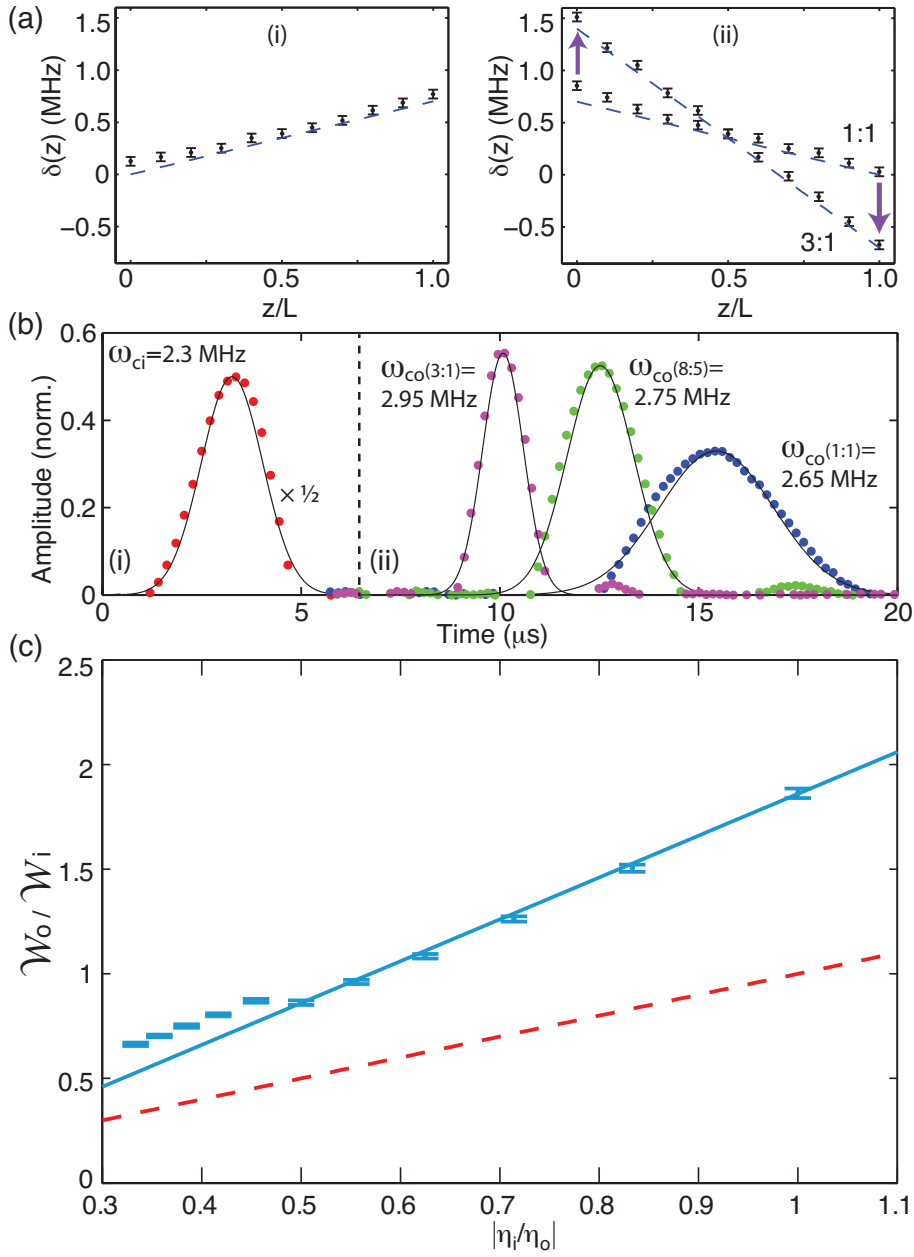


Figure 7.6: Input/Output Bandwidth Characterisation. (a) Two-photon detuning δ as a function of position z along the memory (normalised to length L) due to (i) input gradient and (ii) minimum and maximum output gradients (ratios noted on figure). Blue (dashed) line corresponds to the desired field and points correspond to the measured magnetic field (error bars due to sensitivity of Gauss-meter). (b) Amplitude plot, normalised to size of input pulse, showing (i) input pulse (shown in red, scaled by a factor of $1/2$), (ii) output pulses recalled with varying output gradients as noted. Points correspond to demodulated data, lines correspond to Gaussian fit to data, and ω_c values correspond to the centre frequencies of pulses relative to the LO. Bracketed ratios indicate η_o/η_i . (c) The FWHM of the output pulses \mathcal{W}_o normalised to the FWHM of the input pulse \mathcal{W}_i , as a function of input gradient over output gradient $|\eta_i/\eta_o|$. Points represent measured FWHM (error bars from standard deviation of 100 traces), red (dashed) line corresponds to Eq. 7.4, blue (solid) line corresponds to a numerical simulation, with free parameters: $g = .066 \text{ s}^{-1}$; $\gamma_o = 0$; and $gN/c = 1000$.

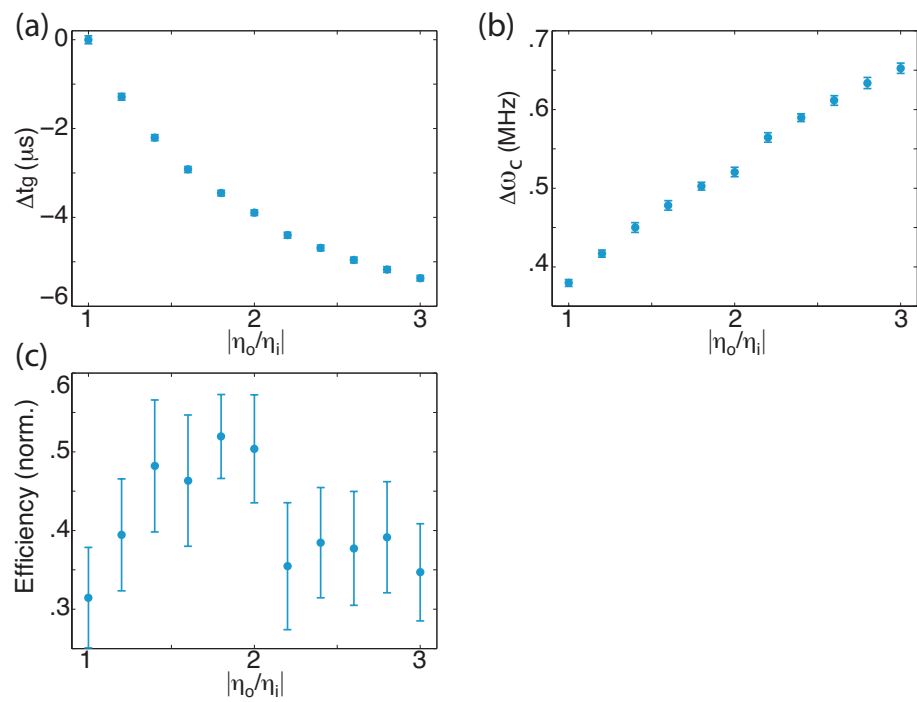


Figure 7.7: Extra Input/Output Bandwidth Characterisation. (a) Change in recall time Δt_g as a function of output to input gradient ratio $|\eta_o/\eta_i|$ compared to recall with $|\eta_o/\eta_i| = 1$. (b) Change in centre frequency of the output pulses relative to the input pulse $\Delta\omega_c$ as a function of $|\eta_o/\eta_i|$. (c) Efficiency as a function of $|\eta_o/\eta_i|$. In all figures points represent data and error bars are from standard deviation of 100 traces.

and memory bandwidth parameters, also showed pulse broadening in excellent agreement with the experimental data.

It can also be seen from Fig. 7.6(b) that the echoes were emitted from the memory earlier, i.e., at a time $t < 2\tau$, when recalled with a steeper output gradient. This was because a steeper gradient will cause the rephasing process to occur at a faster rate, and will also reduce the amount of dispersion. A characterisation of this effect is shown in Fig. 7.7(a). It was also found that the frequency of the output echo was not the same as the input. This is characterised in Fig. 7.7(b). As can be seen, there was nearly a 300 kHz change in ω_c . Though it is possible that a fraction of this was due to an offset in the magnetic field, Fig. 7.6(a)(ii) shows that this offset was smaller than 300 kHz. Instead, most of the change in frequency of the echo can be attributed to the inherent GEM frequency shift predicted in Ref. [224] (see Sec. 3.4.2), which is greater for shorter storage times. Finally, the efficiency as a function of recall gradient is shown in Fig. 7.7(c). This shows that there was an optimal output to input gradient ratio for efficient storage around $\eta_o/\eta_i = 1.5$. This is due to the competing effects of having the pulse leave the memory earlier, therefore being subject to less decoherence inside the memory, and the loss of recall efficiency caused by recalling with a steeper gradient, as shown in Eq. 3.54.

7.3.3 Spectral Filtering

If we now consider the storage of a modulated pulse, the frequency-encoding nature of GEM will mean that the carrier and sideband components of the pulse will be stored in different parts of the memory. Therefore, if we had fine enough control over the recall gradient, we could choose when to recall the different frequency components by switching the gradient only in the pertinent part of the memory. The finer the control over the gradient, the greater the number of frequency components we can individually address.

An experimental demonstration of this filtering is shown in Figs 7.8(a)-(b). Here a carrier pulse with a Gaussian envelope and two frequency components, separated by 700 kHz (the approximate bandwidth of the Gaussian), is sent into the memory. This multi-frequency pulse was created by mixing two sine waves, one at the carrier frequency $\omega_o = 79.65$ MHz and one at the sideband frequency $\omega_{sb} = 350$ kHz. This created two sine waves at $\omega_c \pm \omega_{sb}$, and this signal was then sent to the probe AOM to create a two-frequency-component Gaussian. By reversing the gradient only in one half of the memory at a time, the different frequency components of the pulse could be recalled separately. The output pulses in this case both had the same ω_{co} due to the offset in the recall gradient for the lower frequency component.

Furthermore, following the same logic, by being able to switch the gradient slowly along the length of the memory, the stored pulse can be recalled as its Fourier transform. This is shown experimentally in Figs 7.8(c)-(d) for a modulated Gaussian with two sidebands at ± 700 kHz. In this case the gradient was reversed in three stages, rather than a gradual reversal of the entire memory as was suggested in Ref. [185], due to the limitations of the LabVIEW code refresh rate ($\approx 1 \mu\text{s}$). The three outputs were fit separately on account of their different frequencies. The time window for each demodulation is denoted by the dashed lines in Fig. 7.8(d). As the different frequency components were located in different parts of the memory, we could also choose to recall only certain frequency components by only reversing the gradients in the pertinent parts.

In the above three-frequency-component experiment, the pulse was made in a similar

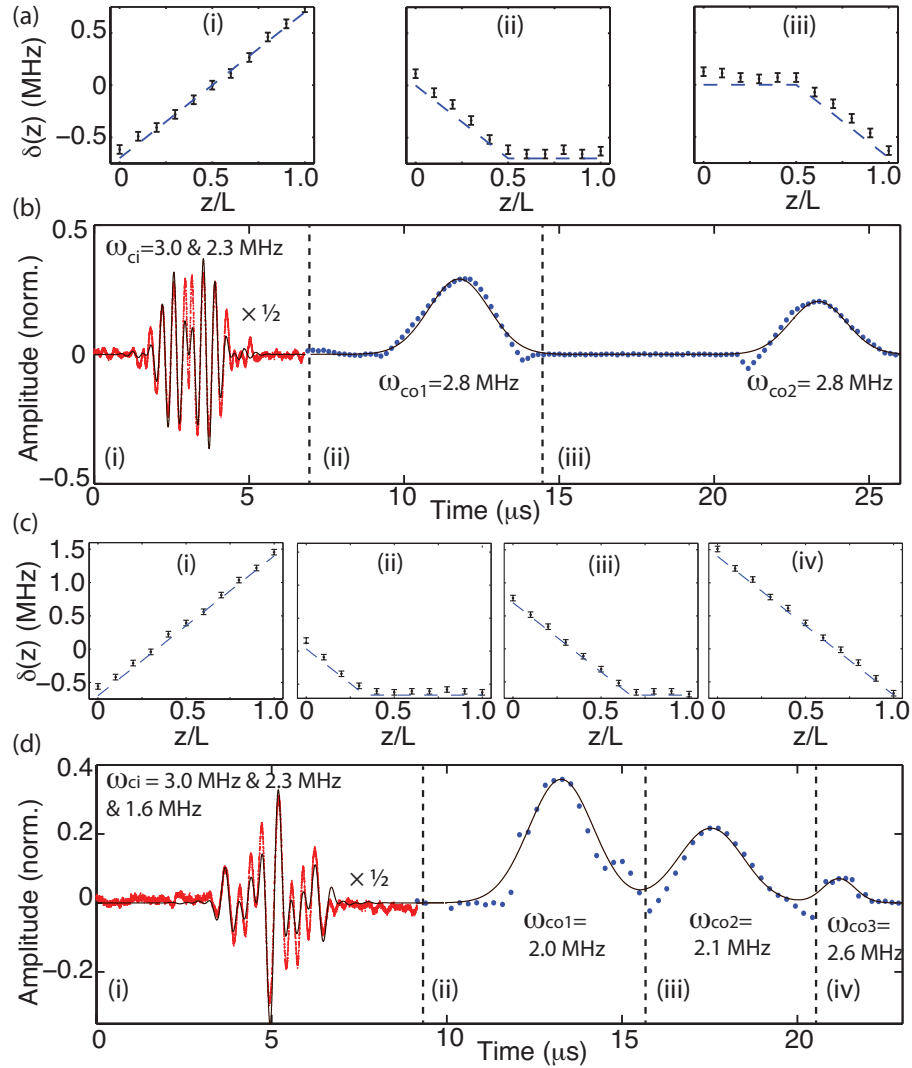


Figure 7.8: Spectral Filtering. (a) Two-photon detuning δ as a function of position z along the memory (normalised to length L) due to gradients (i)-(iii) corresponding to times (i)-(iii) in (b). For traces (a) and (c) blue (dashed) lines correspond to the desired field and points correspond to the measured magnetic field (error bars due to sensitivity of Gauss-meter). (b) Spectral filtering of (i) a Gaussian envelope containing two frequency components separated by 700 kHz (red, non-demodulated, scaled by $1/2$), and the demodulated retrieval (blue) of (ii) higher, and (iii) lower frequency components averaged over 100 traces. For traces (b) and (d) points correspond to data, lines correspond to fit to data, and ω_c values correspond to centre frequencies of pulses. (c) Two-photon detuning due to (i) input and (ii), (iii), and (iv) output gradients corresponding to times (i)-(iv) in (d), which shows the conversion from the time- to frequency-domain of (i) a Gaussian pulse with two modulation sidebands at ± 700 kHz (red, non-demodulated, scaled by $1/2$), and the demodulated retrieval of (ii) higher frequency sideband, (iii) carrier, and (iv) lower frequency sideband averaged over 100 traces (blue).

method to the previous, two-frequency-component pulse: the carrier frequency was mixed with a signal at ω_{sb} , and then this signal was mixed with another signal at ω_{sb} . By altering the relative phase of the two ω_{sb} frequency sources we were able to create a carrier and two sidebands with equal power. The reason this was not the case for the echo was because of coupling field-induced scattering (see Sec. 5.3.4). In normal Λ -GEM storage the coupling field is switched off during the storage process to limit this effect. This was not possible, however, for the multi-pulse recall in a single gas cell as the coupling field must be present for recall to occur.

7.3.4 Pulse Interference

A time-reversal of the spectral filtering process is also possible. That is, if we take two pulses with different frequencies and store them one-at-a-time in different halves of the memory we can alter the gradients in the different halves at different times. This will cause the recalled echoes to overlap, and therefore interfere, at the output of the memory. Previous experiments in pulse interference using GEM have shown how the memory can facilitate interference between modes separated in either the time- or frequency-domains [10]. In that case, pulses were used to change the absorption and recall of other pulses since the light was being stored in the same region of the atomic ensemble. In these experiments we used the MEC to separate the storage regions in space, allowing independent storage and recall of the modes.

Different-frequency pulse interference is presented in Figs 7.9(a)-(b). Here, two pulses separated in frequency by 700 kHz were stored in separate halves of the memory. The lower (higher) frequency pulse being stored in the second (first) half. Setting the gradient to zero in the second half of the memory when the first (lower frequency) pulse P_1 enters ensured it was stored in the first half of the memory. Setting the gradient to zero in the first half of the memory while the higher frequency pulse P_2 entered and was stored served two purposes: apart from ensuring that none of P_2 was stored in the first half of the memory, it also meant that the stored P_1 would not undergo any additional dephasing. Therefore, after P_2 was stored, we could reverse the gradient across the entire memory at once, causing the superposition of the echoes on the output.

The interference pattern will be of the form

$$A_1 \text{Cos}[2\pi\omega_{o1}(t - t_o)] \cdot \text{Exp} \left[\frac{-(t - t_{g1})^2}{2\sigma_{t1}^2} \right] - A_2 \text{Cos}[2\pi\omega_{o2}(t - t_o) + \phi] \cdot \text{Exp} \left[\frac{-(t - t_{g2})^2}{2\sigma_{t2}^2} \right]. \quad (7.5)$$

The terms $A_{1/2}$, $\omega_{o1/2}$, etc. are the same as from Eq. 7.3 except for ϕ , which is the relative phase between the two pulses. The numerical subscripts denote which coefficients belong to which echo (i.e., 1 for E_1 and 2 for E_2). To investigate the phase preserving quality of the memory, we altered ϕ and looked at the phase of the interference pattern of the echo. As can be seen from Fig. 7.9(c), the change in relative phase of the two input pulses matched the relative phase of the interference pattern at the output. The only free parameter in the fitting of the echoes was the relative phase, with the amplitudes, timings and frequencies of the two individual pulses determined from fittings to the storage of individual echoes E_1 and E_2 , shown in Fig. 7.9(b) inset. These parameters are shown in Tab. 7.1.

Two pulses with the same frequency can also be interfered in a similar manner. An experimental demonstration of same-frequency interference is shown in Fig. 7.10. The two

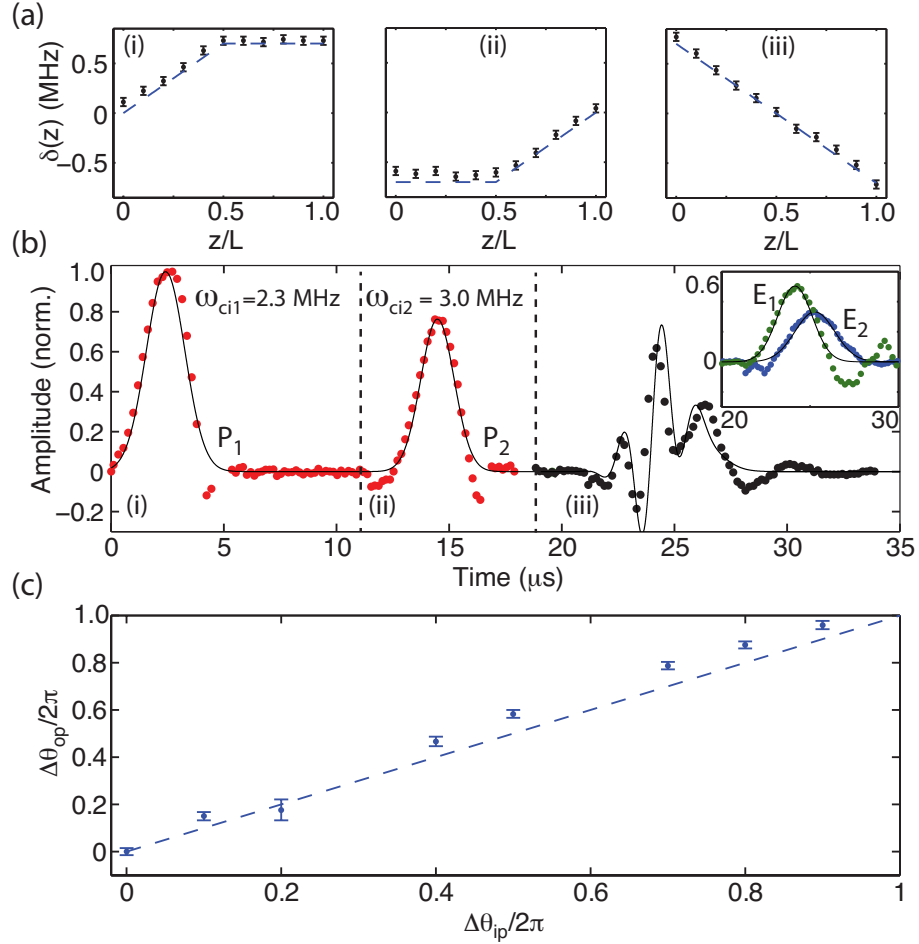


Figure 7.9: Interference with Pulses of Different Frequencies. (a) Two-photon detuning δ as a function of position z along the memory (normalised to length L) due to (i)-(ii) input gradients and (iii) output gradient corresponding to times (i)-(iii) in (b). Blue (dashed) lines correspond to the desired field and points correspond to the measured magnetic field (error bars due to sensitivity of Gauss-meter). (b) Interference of two initially time-separated pulses (i) P_1 and (ii) P_2 , shown in red, which are also separated in frequency by 700 kHz. (iii) Shows the superposition of the two pulses. The inset shows the output from the memory for storage of only a single pulse: P_1 recall (E_1 , green); or P_2 recall (E_2 , blue). Points correspond to demodulated data averaged over 100 traces, lines correspond to Gaussian fit to data, and ω_c values correspond to centre frequencies of pulses. (c) The change in relative phase of the fitted interference pulse $\Delta\theta_{op}$ as a function of the relative phase of the input pulses $\Delta\theta_{ip}$. Points represent data extracted from fit (error bars from standard deviation of 100 traces), and the dashed line corresponds to the theoretical behaviour.

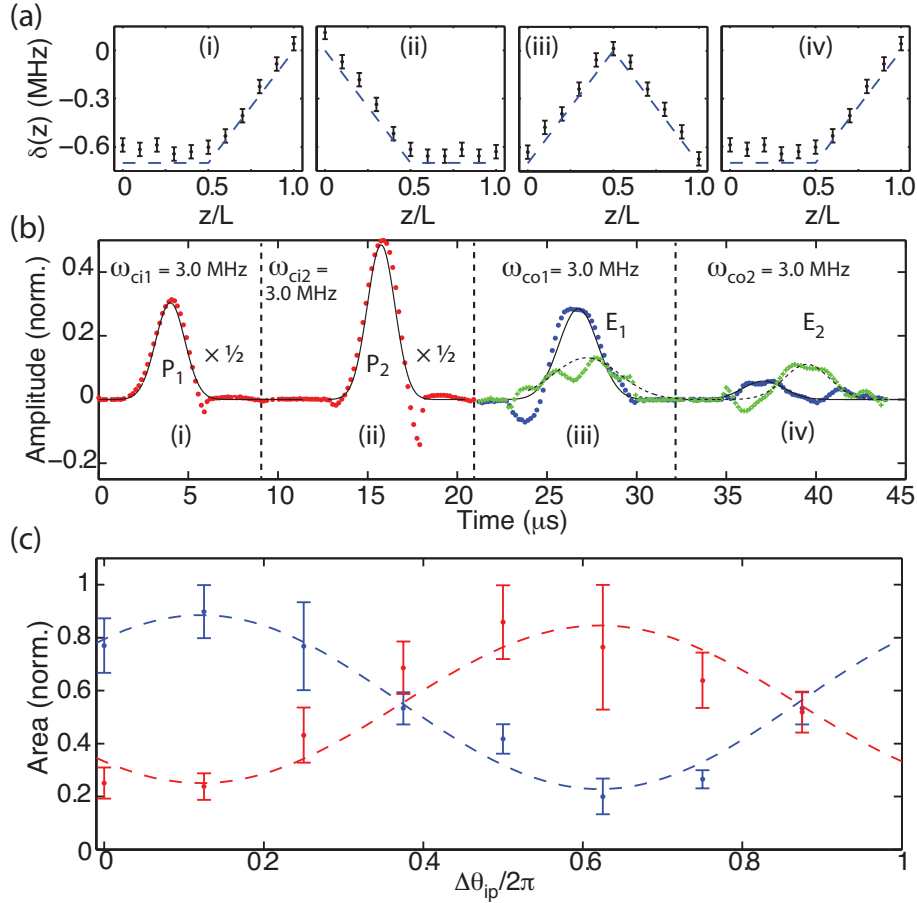


Figure 7.10: Interference with pulses of the same frequency. (a) Two-photon detuning δ as a function of position z along the memory (normalised to length L) due to (i)-(ii) input gradients and (iii)-(iv) output gradients corresponding to times (i)-(iv) in (b). Blue (dashed) lines correspond to the desired field and points correspond to the measured magnetic field (error bars due to sensitivity of Gauss-meter). (b) Interference of two initially time-separated pulses (i) P_1 and (ii) P_2 (red dashed lines), which have the same centre frequency. (iii) Initial E_1 , and (iv) secondary E_2 superpositions of the two pulses. Blue points and solid line correspond to maximum constructive interference for E_1 , while green crosses and dashed line correspond to maximum constructive interference for E_2 . Points correspond to demodulated data averaged over 100 traces, lines correspond to Gaussian fit to data, and ω_c values correspond to centre frequencies of pulses. (c) The change in area (normalised to the maximum intensity of the individual echoes) as a function of the relative phase of the input pulses $\Delta\theta_{ip}$ for (i) E_1 , and (ii) E_2 . Points represent data extracted from fit (error bars from standard deviation of 100 fits), and the dashed line corresponds to a fit to the data.

Parameter	E_1	E_2
A (norm.)	0.84	1
ω_o (MHz)	2.79	2.23
t_g (μ s)	1.125	0
σ (norm.)	1.15	1

Table 7.1: Interference Echo Parameters. The fit parameters for the individual echoes E_1 and E_2 shown in the inset of Fig. 7.9. A and σ values were normalised to E_2 , ω_c was measured relative to the LO frequency, and t_g was determined relative to E_2 .

pulses were, as before, stored in different halves of the memory by setting $\eta = 0$ in the opposite half. As can be seen from Fig. 7.10(a)(iii), the first recall gradient was no longer monotonic and, therefore, the pulse stored in the first half of the memory P_2 was partly reabsorbed in the second half. This is why P_2 had a greater amplitude than P_1 . The non-absorbed component of P_2 interfered with the retrieved light from the second part of the memory. If these two echoes were in phase then there was constructive interference and an enhanced echo E_1 was retrieved. If they were out of phase then E_1 would be reduced in size and the residual energy remained as atomic excitation inside the memory. Therefore, if the gradient in the second half of the memory was switched again, a second echo E_2 could be recalled from the memory from these leftover excitations. This is shown in Figs 7.10(b)(iii)-(iv). These results are similar to those presented in Ref. [10] where another technique was used to cause interference of pulses using GEM.

Figure 7.10(c) shows the areas of the two echoes as the relative phase between P_1 and P_2 was varied. Interference fringes can be seen, with a π phase shift between the two echoes. The visibility for both E_1 and E_2 was approximately 60% (normalised to the maximum echo output for E_1 and E_2 separately). These values could potentially be improved with finer gradient control, especially in terms of timing. Table 7.1 gives an indication of the mis-match between the two echoes retrieved from the memory.

7.4 Cross-Phase Modulation

To further demonstrate the flexibility of GEM when it comes to pulse manipulation, and improve its utility for quantum computing applications, we present here a proof-of-principle experiment where a signal field was used to controllably change the phase of a stored probe pulse via cross-phase modulation (XPM) using the optical Kerr effect. The ability to produce a large XPM would allow GEM to be used for Fredkin gates [295], quantum-phase gates [296], as well as the possibility of using the scheme to generate highly-entangled multi-qubit states, known as cluster states, which are useful for one-way quantum computing [69, 70, 71].

Being able to perform these operation using a memory could prove advantageous as, stored inside a memory, the two signals can be made to interact for long periods of time. To date, experiments on XPM in memories have focused on double-EIT systems via the ac Stark effect [335]. These experiments include the work of B.-W. Shiao et al. who used a cold atomic ensemble of Cs atoms to demonstrate a single-photon phase shift of 10^{-6} rad [336]. This was improved upon by H.-Y. Lo et al. who used a cold atomic ensemble of

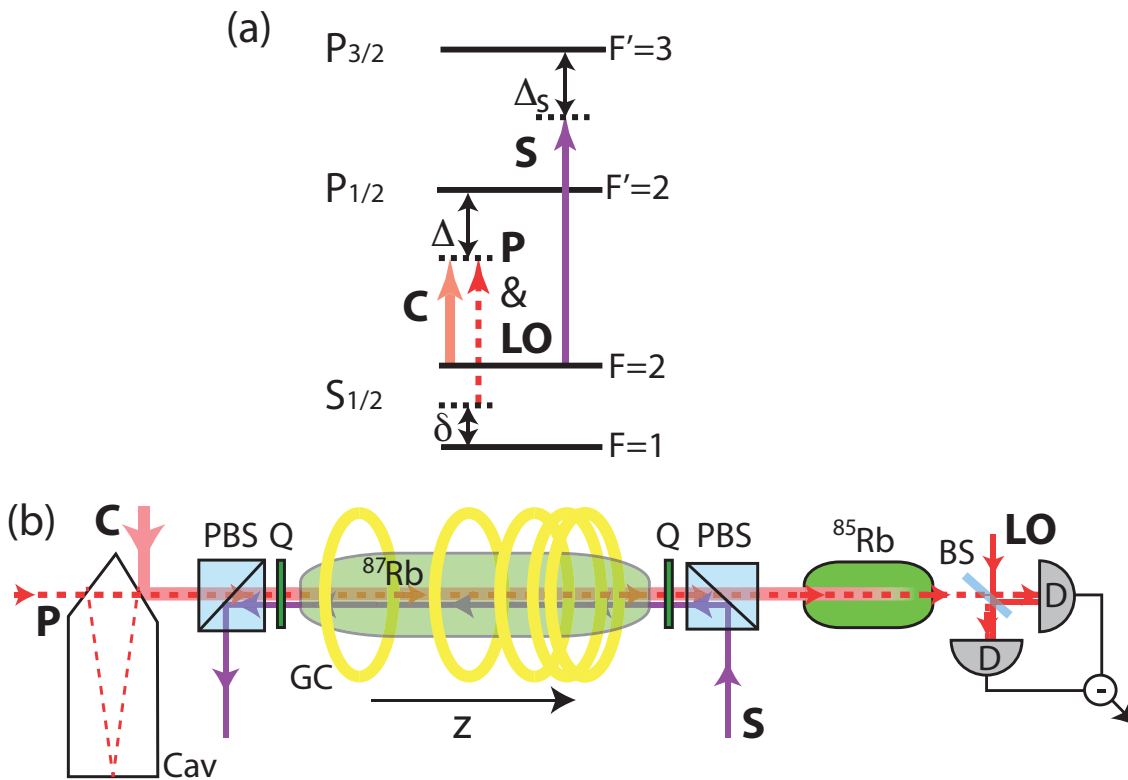


Figure 7.11: Cross-Phase Modulation Set-Up. (a) Energy level diagram for all fields present: C - coupling; P - probe; LO - local oscillator; S - signal. Also shown are one- and two-photon detunings for probe and coupling fields (Δ and δ respectively) and detuning for signal field Δ_s . (b) Schematic diagram of set-up: Cav - ring cavity; PBS - polarising beam-splitter; Q - quarter wave-plate; GC - GEM variable pitch solenoid coils; ^{87}Rb - isotopically pure memory vapour cell; ^{85}Rb - natural mixture vapour cell for filtering the coupling field; BS - 50:50 non-polarising beam-splitter; D - detector.

^{87}Rb atoms to produce an inferred single-photon phase shift of $1.4 \cdot 10^{-5}$ rad [337]. XPM has also been used to manipulate images stored in a cloud of ^{87}Rb atoms, again using EIT [338]. However, an approximate analytical solution of the XPM EIT scheme has been derived that shows large loss will occur for large phase shifts, limiting noiseless XPM for EIT-based systems to “*essentially useless values*” [339].

The set-up used to demonstrate XPM using GEM is shown in Fig. 7.11, and is similar to that described in Sec. 6.2, with the MEC being replaced with the two variable-pitch solenoid coils originally used for GEM operation. To this set-up we added a signal field, detuned by approximately $\Delta_s = -2$ GHz from the D2 $F = 2 \rightarrow F' = 3$ transition, using a pair of polarising beam-splitters placed before and after the memory. This signal beam was sent counter-propagating to the coupling and probe beams to avoid contaminating the heterodyne detection. The theoretical phase shift for such a system is given by [12]

$$\phi_{xpm} = \frac{g^2 |\mathcal{E}_s|^2 \Delta_s t_s}{2(\gamma^2 + \Delta_s^2)}, \quad (7.6)$$

where g is the atom-light coupling strength, t_s is the length of the pulse, \mathcal{E}_s is the amplitude of the signal field, and γ is the excited-state decay rate.

An example of the results obtained using this set-up is shown in Fig. 7.12(a). To be able to accurately measure the phase of the echo, as the probe and heterodyne fields were not phase locked, a reference pulse (R) was sent in before the input pulse to provide a phase reference. The coupling field was turned on after this pulse and the input was sent into the memory, where it was stored for approximately 15 μs . During this storage time a 10- μs -long signal pulse could be sent through the memory. The two traces shown in Fig. 7.12(a) were separated by 60 μs , the length of one run. As can be seen from the insert, there was approximately a $\pi/2$ phase shift between the two pulses caused by the signal field. There was also a reduction in efficiency from 53% to 7% with the signal field present, due to the signal-field-induced scattering.

By varying the power in the signal field from zero up to approximately 10 mW and measuring ϕ_{xpm} we were able to demonstrate that the behaviour of the system follows Eq. 7.6. This is shown in Fig. 7.12(b). From these results we were able to determine a phase shift on the order of 10^{-12} rad for single-photons in the probe and signal beams. We also showed that there is no self-phase modulation present by varying the power in the probe beam \mathcal{E}_p for a variety of different signal field powers. This is shown in Fig. 7.12(c).

7.5 Discussion

In this chapter we have experimentally demonstrated a number of different spectral manipulation operations using GEM. These operations could have various uses in a quantum information network. For instance, the ability to alter the bandwidth of the pulse, demonstrated in Sec. 7.3.2, could be used to match systems with different bandwidths and, by increasing bandwidths, help to improve bit rates for a number of time-bin qubits as described in Ref. [293]. Combining this with the ability to change the centre frequency of the stored information, demonstrated in Sec. 7.3.1, would allow one, in theory, to match any two optical systems.

Being able to controllably alter the frequency of stored information would allow for the conversion of time-bin qubits into frequency-bin qubits. It would also, along with

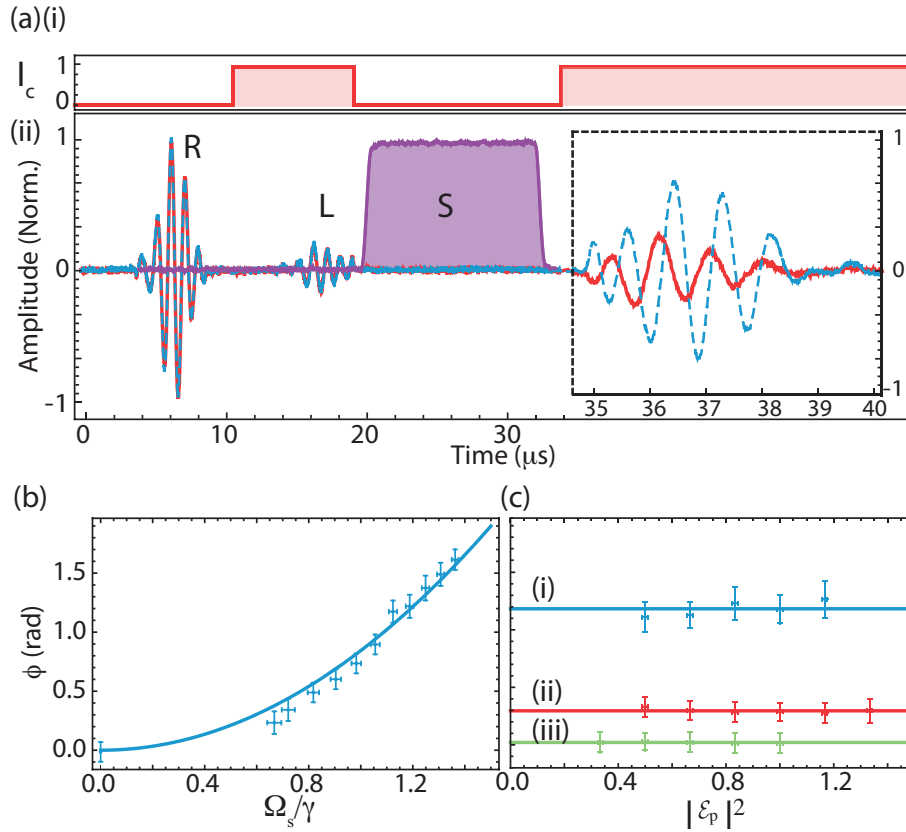


Figure 7.12: Cross-Phase Modulation Results. (a) (i) Coupling field intensity I_c during run, and (ii) heterodyne data for one run with the signal pulse (red, solid line) and one without (blue, dashed line), taken $60 \mu\text{s}$ apart. R - reference pulse; L - probe field leakage; S - signal pulse, inset shows a close-up of the echo recall. (b) Echo phase shift as a function of signal field Rabi frequency $\Omega_s = g |\mathcal{E}_s|$, normalised to the decay rate γ . Points represent experimental results (error bars from standard deviation of 1000 traces) and the solid line corresponds to a theoretical fit using Eq. 7.6. (c) Echo phase shift as a function of probe field intensity $|\mathcal{E}_p|^2$ for different values of signal field Rabi frequency: (i) $\Omega_s = 1.12\gamma$; (ii) $\Omega_s = .72\gamma$; (iii) $\Omega_s = 0$. Points represent experimental results (error bars from standard deviation of 1000 traces), and the solid line corresponds to expected values if there is no self-phase shift. Adapted from Ref. [340].




Figure 7.13: Memory Network. Schematic diagram of potential network of memories with: C_x - coupling field number x ; QM_x - memory number x ; PBS - polarising beam-splitter.

the frequency encoding nature of GEM, allow for the ability of frequency multiplexing. A number of pulses with different frequencies could be combined into one temporal pulse inside the memory, as demonstrated in Sec. 7.3.4, and sent down the communication channel. Once they reached the other end of the channel they could be separated with a second memory, as demonstrated in Sec. 7.3.3. This could greatly improve qubit rates over optical channels in quantum information networks, as well as improve quantum repeater protocols as discussed in Sec. 7.1. It is worth mentioning that this multiplexing is only possible with a memory that has the ability to both combine and separate multiple frequency components inside the memory into/from one temporal pulse. This ability is made possible by the precision control over the magnetic field provided by the MEC. This is not possible with the polaritonic interference techniques described in Ref. [10].

The phase-sensitive interference of initially time-separated pulses, demonstrated in Sec. 7.3.4, could also find applications in quantum computing. In an all optical switch, for instance, where it is the relative phase between the pulses that determines how much light is emitted at different times [341].

The XPM experiment, shown in Sec. 7.4, could also find use in quantum computing architectures, though a much higher phase shift would be desirable. The single-photon phase shift of 10^{-12} rad could be improved by moving the signal field closer to resonance. The reason we did not reduce the detuning in the experiment was that it would lead to even greater loss of efficiency through signal-field-induced scattering. Moving to cold atomic ensembles, as used in the EIT XPM experiments [337, 336, 12], should allow us to reduce this detuning by orders of magnitude as the low temperatures will greatly reduce the Doppler shift of the atoms. The smaller interaction volumes of cold atomic ensembles could also lead to six orders of magnitude improvement in phase shift. Finally, storing the signal field, as opposed to simply passing it through the ensemble, could also improve the phase shift by another two orders of magnitude. Altogether, these improvements could lead to an optimal predicted single photon phase shift of 10 mrad [258], orders of magnitude larger than what is possible in EIT systems [337, 336].

All these potential applications require high efficiencies and therefore high optical depths. A high optical depth is especially important as increasing the bandwidth of the system will decrease the recall efficiency. The pulse width for the MEC experiments was chosen to match the memory bandwidth to optimise efficiency. However, it was seen in Sec.s 7.3.1 and 7.3.2 that this led to dispersion of the echoes. Being able to increase the

memory bandwidth enough such that the pulse is absorbed away from the dispersive edges, as shown in modelling in Ref. [252], would significantly reduce this effect. If an increase in bandwidth were not possible, then a comprehensive characterisation of the effect, like the one presented in Sec. 7.3.2 would be required to make this operation useful.

A drawback to using one physical memory (i.e., the 20-cm-long vapour cell) as a system of n sub-memories is that the optical depth for each individual sub-memory will be $1/n$ of the total memory optical depth. An alternative method would be to use n physical memories placed in series to create a memory network. Not only would this increase the overall optical depth of the system but it could help to alleviate two other drawbacks to the sub-memory approach taken here: finer control of the gradient; and coupling-field-induced scattering.

Much care was taken with the design and construction of the MEC to allow for maximum precision in gradient creation with realistic currents. The MEC allowed for separate control over the magnetic field offset and bandwidth down to tens of kilohertz, something that was not possible with previous coil designs. This allowed us to characterise the bandwidth and frequency manipulation properties of the memory. It also allowed for individual storage and recall of up to three frequency components from the memory, limited by the resolution of the MEC. Using many physical memories as one memory network would automatically increase this resolution with respect to the length of each sub-memory. Another option for improving the resolution would be to move to an alternate gradient creation technique such as the ac Stark effect (see Ch. 10), which could allow for a much sharper change in the gradient by using, for instance, a spatial light modulator to create the intensity profile of the gradient field.

The coupling-field-induced scattering was discussed in Sec. 7.3.3 with regards to the extra decay of information that is left in the memory while other information is recalled. This is a concern as this scattering leads to a large increase in the decoherence rate (see Tab. 5.1) and cannot be combatted in a single memory. However, this issue could be addressed with a network of quantum memories $QM_1 \rightarrow QM_n$ if one were to use orthogonal polarisations for the probe and coupling fields and place polarising beam-splitters between the memories. This is illustrated in Fig. 7.13.

7.6 Conclusions

In this chapter we have presented experimental demonstrations of spectral manipulation operations theoretically proposed in Ref. [185]. We showed that using the gradient echo memory scheme we can alter the bandwidth (and therefore temporal profile) of a pulse, as well as change its centre frequency. We also demonstrated the ability of GEM to act as a spectral filter and, using the frequency-encoding nature of GEM, were able to recall a modulated pulse as its Fourier transform. We showed that two initially time-separated pulses, with the same or different frequencies, could be caused to interfere coherently at the output of the memory. Finally, in a separate experiment, we presented a proof-of-principle demonstration of cross-phase modulation in GEM, obtaining a single-photon phase shift of 10^{-12} rad with no self-phase modulation. These abilities could be used to improve qubit rates across quantum communication channels, as well as potential uses in quantum computing applications.

Spatially Multi-Mode Storage Using GEM

Mozzie: “You know that art studio where you keep all the paintings you’ve done?”

Neil: “Yes, what about it?”

Mozzie: “Now it’s just a studio.”

- “On Guard” *White Collar*, USA Network, 7 June 2011, TV Series.

8.1 Introduction

In the previous chapter we discussed how having a multi-mode quantum memory can lead to improved quantum repeater (QR) designs by increasing bit rates by orders of magnitude [105] and reducing their sensitivity to memory coherence times [106]. We have also experimentally shown in previous chapters both the temporal (Ch. 5) and frequency (Ch. 7) multi-mode nature of GEM. Here we investigate another type of multi-mode storage: spatially multi-mode storage.

Apart from the above-mentioned QR improvements, being able to store different spatial modes can lead to orders of magnitude improvement in storage capacity. A spatially multi-mode memory could also potentially find use in novel quantum information architectures such as quantum holographic teleportation [342, 343], quantum dense image coding [344] or holographic quantum computing [345]. Other applications that make use of information embedded in images, such as quantum secure imaging [346], could also potentially make use of a memory that can store these states.

It is therefore no surprise that much experimental work has already been carried out towards developing a spatially multi-mode quantum memory. In warm vapours, electromagnetically-induced transparency (EIT) has been used to store images for up to $30 \mu\text{s}$ [263, 347, 262], with a diffusion constant of $D = 10 \text{ cm}^2 \text{ s}^{-1}$. By using elegant Hermite-Gauss (HG) modes [348] the group of D. Yankelev et al. were recently able to demonstrate storage for up to $80 \mu\text{s}$ [349], still using EIT. Delay of entangled images has also been demonstrated using four-wave mixing (FWM) in warm vapour [194].

The low atomic speeds in cold atomic ensembles have allowed not only image storage, but multiplexed image storage [327, 350, 338]. In the last of these papers, the group of L. Zhao et al. also demonstrated manipulation of the image using cross-phase modulation. Storing images using FWM has also been demonstrated with cold atoms [351]. Reference [347] used FWM in their EIT storage process.

Solid state systems have no atomic motion at all and have therefore led to the longest storage time for image storage of 1.2 ms, again using EIT [352]. As with other long

storage times in solid state ensembles, RF pulses were used to improve coherence times. Using the same set-up, the authors were also able to demonstrate angular and frequency multiplexing, as well as a combination of the two [326].

In this chapter we demonstrate spatial multi-mode storage using Λ -GEM in a warm vapour cell. In Sec. 8.2 we describe the set-up used to detect different spatial modes, before presenting results showing a conjugate FWM pulse in Sec. 8.3 and different HG mode storage in Sec. 8.4. We also show that we can selectively store and recall a particular maxima of an HG mode. This is followed by storage of images in Sec. 8.6.

At the time these experiments was being undertaken, another group was carrying out similar work. Q. Glorieux et al. showed that, using GEM in a ^{85}Rb vapour cell, they could store two images, one after the other, and then recall them, with a maximum efficiency of 8% and a maximum storage time of $4.5 \mu\text{s}$ [353].

Most of the work presented in this chapter has been published in the journal article:

- D. Higginbottom, B. M. Sparkes, M. Rancic, O. Pinel, M. Hosseini, P. K. Lam, and B. C. Buchler, *Spatial-mode storage in a gradient-echo memory*, Physical Review A **86**, 023801 (2012).

8.2 Spatial Mode Detection Set-Up

The set-up for this experiment, shown in Fig. 8.1(a), was very similar to that used in the previous two chapters and, as such, will not be discussed in detail here. However, the sizes of the probe and coupling beams were different, with waists of 1.5 mm and 3 mm respectively. The multi-element coils from the previous chapter were also replaced with the variable-pitch solenoids from Ch. 6.

The main difference from previous set-ups was the detection scheme. As we needed to be able to detect multiple spatial modes we could no longer easily use heterodyne detection. Instead we used a single photo-diode, which had no mode selectivity, for measuring the power in the probe beam. To be able to image the various modes we also used triggerable charged-coupled device (CCD) cameras (Point Grey Grasshopper2) and a flip mirror to move between intensity detection and the CCD image.

A ring cavity was again used to combine the probe and coupling fields. Different modes were selected for storage by maximising, and then locking to, higher order spatial modes. Detectors D1 and D2 were used to find common resonant modes between the locking and probe fields, while CCD1 was used to determine which probe mode was being passed through the cavity. Once a mode was selected, and the cavity aligned, FM1 was removed. Then, depending on whether intensity or image data was required, FM2 was either removed or put in place.

One issue we faced with imaging the probe was that, as the probe beam power was less than one microwatt with a duration of approximately one microsecond, the signal had to be averaged over many traces to obtain a satisfactory signal-to-noise ratio for the CCD measurements. Also, without mode-sensitive detection, the coupling field leakage from the filter cell was detected on the CCD along with the probe. As well as this, the CCDs had a minimum digital shutter time of $30 \mu\text{s}$.

To overcome these issues we altered the LabVIEW triggering code introduced in Ch. 6. Firstly, we set the total run length to $120 \mu\text{s}$, leaving enough time for the CCD to stop taking data before the next input. Then 50-100 traces were accumulated on the CCD (i.e.,




Figure 8.1: Spatial Mode Detection Set-Up. (a) Schematic diagram of set-up: P - probe field; C - coupling field; L - locking field; BS - 50:50 non-polarising beam-splitter; FM - flip mirror; D - detector; CCD - charged-coupled device camera; GC - GEM variable pitch solenoid coils; ^{87}Rb - isotopically pure vapour cell; ^{85}Rb - natural mixture vapour cell for filtering the coupling field. (b) An example batch run protocol: (i) probe only batch, N_i times; (ii) coupling only batch, N_{ii} times; (iii) probe and coupling batch, N_{iii} times; (iv) coupling only batch again, then the whole sequence is repeated. Rectangles indicate one run, $t_w = 4$ ms is the wait time between each batch. All N s between 50-100.




Figure 8.2: Four-Wave Mixing Pulse. Processed CCD image showing (i) conjugate four-wave mixing pulse and (ii) echo.

total echo acquisition time of 6-12 ms) before being sent to the computer as one batch image. This process took 40 ms. After this we took 50-100 runs of just the coupling field leakage to produce another batch image, and then continued to alternate between the two batches to produce approximately 100 images in total. The changes to the LabVIEW triggering code allowed us to have multiple separate runs (i.e., only input, no input, memory, etc.), have them repeat a set number of times and then wait for a user-defined period before moving on to the next batch. This batch process is illustrated in Fig. 8.1(b) and again demonstrates the flexibility LabVIEW provides for data acquisition.

Once all the data was collected, it was processed to remove as much coupling field background, as well as ambient light, as possible. The latter was removed by ensuring the room lights were off during the experiment, as well as using a border of 100 pixels from the outside of the image (i.e., not containing any probe or coupling light) and subtracting the average intensity of this region from the rest of the image. To remove the remain coupling field from the images we took one batch of coupling only images before and after each echo batch. We then averaged the coupling field intensity from both sides and subtracted this from the echo image. This method allowed us to overcome jitter in the beams caused by heating of the air around the vapour cells.

8.3 Four-Wave Mixing Pulse

We discussed the negative effects of four-wave mixing (FWM) on memory storage, in terms of amplification and adding noise, in Sec. 5.3.5. Until this point, using mode- and frequency-sensitive heterodyne detection, we had not observed FWM during normal memory operation. Indeed, the noise measurement results presented in Ch. 6 show that if there was any added noise due to FWM then it had a minimal effect at the low photon-number-levels used.

However, an added benefit of using a new method of detection is that previously undetected effects can be discovered. In this case we found that we could, with the right parameters, produce a conjugate pulse along with the echo. To determine the origin of this pulse, we looked at the effect of varying the coupling beam angle. We found that we could steer the conjugate pulse by changing the coupling beam alignment, while the probe remained stationary (but with reduced efficiency). This behaviour suggested that the conjugate was a result of coherent anti-stokes FWM, discussed in Sec. 5.3.5. This is

because the output wave-vector was seemingly determined by conservation of momentum (Eq. 5.15), coming out at an angle between the probe and coupling beams. We also investigated the effect of varying the two-photon detuning. We found that the energy in the conjugate was optimised when the two-photon detuning was placed at the edge of the Raman absorption line. This is also where, theoretically, the FWM amplification should be greatest (see Sec. 5.3.5). The frequency of the FWM pulse would be 6.8 GHz detuned from the probe frequency and would not be detected with the heterodyne set-up. It is also possible that the more intense probe needed for CCD detection, and the increased possibility of seeding the FWM pulse with carrier light leaking through the first cavity that accompanies this (see Fig. 6.1), led to an increase in the FWM effect. The significance of FWM for memory operation is described in Sec. 5.3.5.

As we no longer had the benefit of heterodyne detection to filter out this pulse, we had to remove it from the data by other means. As FWM is very sensitive to the optical depth we lowered the memory cell temperature from roughly 80°C to 70°C for the results presented in the remaining sections. We also used the coupling beam angle to steer the conjugate pulse out of the probe beam, and altered the other memory parameters (coupling power, magnetic fields etc.) to remove this pulse from the data, sacrificing memory efficiency in the process.

8.4 Spatial Mode Storage Results

Using the set-up presented in Sec. 8.2 and the parameters mentioned in the previous section to avoid the FWM pulse, we were able to investigate the efficiency of different mode storage using GEM. The efficiency in this case was measured by integrating the intensity of the processed CCD images, with the coupling field turned off during the storage window to minimise the decoherence rate. Figures 8.3(a)-(e) show decay traces for the Hermite-Gauss transverse electromagnetic (TEM) modes TEM_{00} , TEM_{01} , TEM_{11} , TEM_{02} and TEM_{12} modes respectively. As can be seen from the insets, GEM preserves the pulse shape nicely over the 12 μs storage time shown (5 μs for the TEM_{12} mode and 12 μs for the TEM_{03} mode).

It can be seen from the fit to the data of the form $\text{Exp}[-t/t_e]$ in Fig. 8.3(a) that the decay of the TEM_{00} mode was exponential. The time constant of this decay was 47 μs . If there was no transverse diffusion effects present in the memory we would expect all modes to decay at the same rate. As can be seen from the comparison of the red to the blue lines in Figs 8.3(b)-(e), this was obviously not the case, and we saw faster decay rates for all higher order modes. In general, the higher the mode order, the faster the decay. For instance, the fitted exponential decay time for the TEM_{01} was 37 μs , TEM_{11} was 23 μs and TEM_{12} was 20 μs . However, the decay for the TEM_{02} was 50 μs and TEM_{03} (decay not pictured) was 42 μs .

This can be explained by diffusion effects: in the TEM_{01} and TEM_{11} modes, 50% of the light is out of phase with the other 50%, unlike the TEM_{00} mode where all light is in phase, and therefore diffusion of atoms between adjacent nodes will result in lower efficiency. This can be seen from Fig. 2.3. This would also explain the TEM_{02} and TEM_{03} modes having a longer coherence time, as they have more than 50% of the power in the outside intensity maxima and therefore diffusion between adjacent maxima will have less of an effect, leading to less loss for the same amount of atomic motion.

Most effects discussed above were modelled numerically using XMDS2 [334], and are presented in Ref. [280].

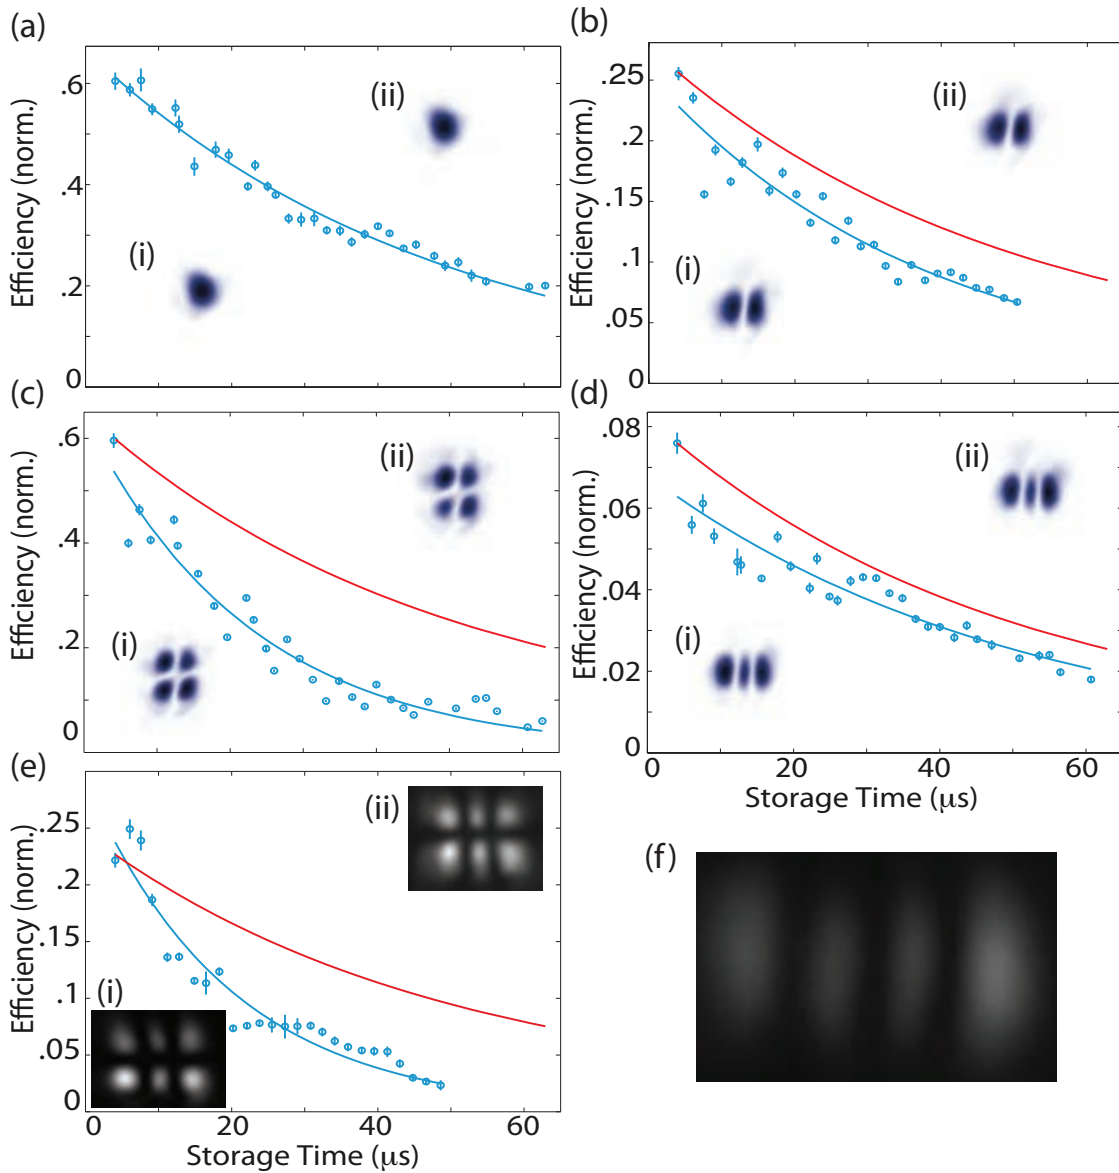


Figure 8.3: Multi-Mode Storage in GEM. Decay curves and CCD images for different modes stored using GEM: (a) TEM₀₀; (b) TEM₀₁; (c) TEM₁₁; (d) TEM₀₂; (e) TEM₁₂. Points correspond to an average of approximately 100 images, error bars are standard deviations, blue lines are exponential fits to the data, (b)-(e) red lines correspond to the decay constant in (a) normalised to highest efficiency. (a)-(d)(i) A processed CCD input image; (a)-(d)(ii) a processed CCD echo image after 12 μs; (e)(i) a non-processed CCD input image; (e)(ii) a non-processed CCD echo image after 5 μs. (f) A non-processed CCD images of a TEM₀₃ mode after storage for 10 μs.




Figure 8.4: Selective Maxima Storage. (a) TEM_{01} input pulse, (b) selective storage and recall of left node only, (c) selective storage and recall of right node only.

8.5 Selective Storage

The coupling field is required for both storage and retrieval in Λ -GEM. This fact has already been put to use to allow for pulse resequencing [9]. It can also be used to allow for selective pulse storage and recall by only shining the coupling field on the desired transverse spatial region of the memory. This is illustrated in Fig. 8.4. Figure 8.4(a) shows a TEM_{01} mode, which was sent into the memory. By aligning the coupling field to overlap with only the left or right maxima we were able to selectively store and recall either one separately in the memory (Figs 8.4(b) and (c)).

8.6 Image Storage

Finally, we investigated the storage of arbitrary images in the memory. To create the image on the probe we simply put a mask in the probe path. Then, to combine the probe and coupling fields we had to remove the ring cavity (as it would not pass an arbitrary superposition of modes) and use a 50:50 non-polarising beam-splitter. We also halved the size of the probe to optimise the image storage. This change in set-up is shown in Fig. 8.5(a) and results for different images are shown in Figs 8.5(b)-(d). As can be seen, these images decay inside the memory much faster than the modes shown in Fig. 8.3. This is not surprising due to the much higher spatial frequencies present in these images, with transverse atomic diffusion now causing a blurring out of the image on a much shorter timescale. The memory efficiencies for these images are also low, around 4% for all images, partly due to the 50% loss in coupling field power.

8.7 Discussion

Having another detection mechanism, in the form of the CCD camera and the batch triggering set-up, provided us with another way of investigating our warm vapour Λ -GEM system. The discovery of the FWM conjugate pulse was an example of the usefulness of probing the system in different ways as the mode-sensitive nature of heterodyne detection will be blind to this effect.




Figure 8.5: Image Storage. (a) Schematic diagram of image storage set-up: P - probe field; C - coupling field; M - mask for probe; BS - 50:50 non-polarising beam-splitter; FM - flip mirror; D - detector; CCD - charged coupled device camera; GC - GEM coils; ^{87}Rb - isotopically pure vapour cell; ^{85}Rb - natural mixture vapour cell. (b)-(d) Storage of different images: Batman logo; Ψ ; and The Australian Research Council Centre of Excellence for Quantum Computation and Communication Technology (CQC²T) logo respectively. (i) input image; (ii) echo image after $4 \mu\text{s}$ storage; (iii) echo image after $8 \mu\text{s}$.

As we were working with a warm vapour, it was no surprise to discover diffusion affecting the results. For the TEM₀₀ mode, the coherence time was 47 μ s. For higher mode storage there are a number of ways to improve coherence times. One would be to use an imaging system to store the Fourier transform of the image in the ensemble. This would make the system much less sensitive to diffusion and improve the image stability. This method has been experimentally demonstrated in Ref.s [338, 347, 351]. Another option would be to use a set of modes that are less affected by diffusion, such as the elegant HG modes [348] discussed in the introduction, to improve the recall visibility. For a diffusion-limited memory another option would be to use a new storage medium with less (or no) atomic motion such as a solid state system or cold atomic ensemble (discussed more in Ch.s 9-12).

As well as lower coherence times, the higher order modes and images were stored with lower efficiency than the TEM₀₀ mode. This will be partly due to the increase in the wave-vector mismatch between the coupling field in the TEM₀₀ mode and the higher order mode of the probe and the increase in two-photon Doppler broadening this will cause (see Eq. 5.4). We ideally want to efficiently store *any* pulse, without a priori knowledge of the spatial profile. One option to reduce the dependence of efficiency on spatial mode would be, again, to move to cold atoms where the atomic motion and therefore Doppler broadening will be greatly reduced.

Finally, the proof-of-principle ability of selective storage and recall could provide intriguing possibilities for further manipulation experiments. By incorporating a spatial light modulator in the coupling beam path we would have the ability to store an image and then recalling only selected parts. This provides us with another method, apart from the frequency-dependent recall presented in Sec. 7.3.3 of the previous chapter, or the temporal-dependent resequencing shown in Fig. 3.3(c) for selective information recall, and could be useful in information processing algorithms, or in multiplexing QR protocols.

8.8 Conclusions

In this chapter we have demonstrated the ability of GEM to store multiple spatial modes. The use of charged-coupled device cameras also allowed us another tool for investigating GEM, with the result being the discovery of a four-wave mixing conjugate pulse. We used the CCD cameras to investigate the decay of different spatial modes, finding the coherence time of the TEM₀₀ mode to be 47 μ s. Other modes decayed at a faster rate due to transverse diffusion. We also briefly demonstrated the ability to selectively store and recall parts of a spatial mode. Finally, we demonstrated storage of images. The higher spatial frequencies of these images meant that they decayed at a faster rate than the TEM modes.

Part III

Cold Atom Experiments

Introduction to Light Storage with Cold Atomic Ensembles

Lord Eddard Stark: “Winter is coming” - Game of Thrones (recurring line), HBO, 2011-Present, TV Series (based on the books by George R. R. Martin).

9.1 Introduction - Why Cool Kids Use Cold Atoms

In the following chapters we will move away from warm vapour cells and towards storing light using cold atomic ensembles. The reason for this change of medium is two-fold: firstly, the high densities that can be achieved with cold atomic ensembles could lead to larger optical depths, and therefore even higher efficiencies; and, secondly, the lower speeds of the colder atoms will result in longer diffusion times and potentially less collisional broadening, leading to longer absolute storage times.

In Ch. 8 we measured a coherence time for our warm vapour experiment of approximately $50 \mu\text{s}$. For use in quantum repeater nodes, one of the main motivations behind developing a quantum memory, we must be able to store information for at least milliseconds, if not seconds (see Sec. 1.3.1). In the case of warm vapour memories, coherence times are usually limited by a combination of atomic diffusion and magnetic field decoherence (see Sec. 5.3). Therefore a potential way of improving storage times is to reduce the atomic motion using slow-moving atoms. Laser cooling provides an efficient way to achieve this, and cold atom memories have been the subject of many experiments aiming for longer memory storage times.

Reducing atomic motion is not the only way to increase storage times - it depends on the storage protocol. For example, it has been demonstrated that long storage times can be obtained in an electromagnetically-induced transparency (EIT)-based warm memory by using high buffer gas pressure [124]. In a GEM scheme, however, high buffer gas pressure increases collisional broadening and absorption of the coupling beam [252]. The effect of transverse diffusion can be minimised by increasing the interaction volume and using anti-relaxation coatings to minimise inelastic wall collisions [275, 276]. Although this technique would reduce transverse diffusion, the longitudinal diffusion would eventually limit coherence times. As discussed in Sec. 5.3.2, this will lead to frequency blurring of the pulse and, therefore, inefficient recall

Lower atomic diffusion rates would also be beneficial for image storage, as discussed in Ch. 8, while the longer storage times and smaller interaction volumes could help improve the single-photon phase shift for cross-phase modulation experiments, as discussed in Sec. 7.4. Also, increasing the coherence time while keeping the pulse length constant

would improve the delay-bandwidth product (DBP) for the system (see Sec. 1.4.5).

This chapter presents the basic theory behind the cooling and trapping of atoms for use in light storage applications in Sec. 9.2, as well as a discussion of decoherence effects in cold atomic systems in Sec. 9.3. A literature review of pertinent previous experiments that made use of cold atoms for light storage, and work towards optimising cold atomic ensembles for light storage, is presented in Sec. 9.4.

9.2 Theory of Atom Trapping and Cooling

In this section we present the basic theory needed to understand the workings of the main cooling and trapping processes used to create cold atomic ensembles for quantum memory applications.

The basics of laser cooling of atoms is based on the Doppler shift. That is, atoms seeing light at a different frequency depending on their speed relative to laser beam (see Sec. 5.3.1). Normally this is considered an undesirable effect. However, it is handy for cooling of atoms as, if the detuning of the laser is chosen wisely, atoms will preferentially absorb light only when travelling towards a laser. As there is no preferential direction for spontaneous emission, this will result in a force on the atom away from the laser. By combining beams in the x -, y - and z -directions, and being clever in choices of polarisations, this can be used to cool atoms in three dimensions, and is the foundation of the magneto-optical trap presented in Sec. 9.2.1. From here, more complex forms of cooling (Sec. 9.2.2) and trapping (Sec. 9.2.3) are presented.

9.2.1 Magneto-Optical Traps

The work-horse for cold atom experiments is the magneto-optical trap (MOT). As its name suggests, a MOT combines the radiation pressure of optical fields with a magnetic field gradient to cool and trap neutral atoms. This process is explained below, focusing firstly on the optical fields and then the magnetic fields.

The main cooling mechanism present in a MOT is dissipative (or absorptive) interaction of light with atoms. This is where an atom will absorb, and then re-emit, photons from an optical field. If we imagine a situation where light is shone onto an atom, momentum will be transferred to the atom in the direction of the light field with each absorption event, but the spontaneous emission will be in all directions. This will lead to an overall force in the direction of propagation of the light field. The strength of this scattering force for light of intensity I interacting with an atom with an excited-state decay rate γ and a corresponding saturation intensity I_{sat} will be given by [204]

$$F_{scatt} = \frac{\hbar k \gamma}{2} \frac{I/I_{sat}}{1 + I/I_{sat} + 4\Delta^2/\gamma^2}, \quad (9.1)$$

where Δ is the detuning of the light from the atomic resonance ω_o , and $k = 2\pi/\lambda$ is the wave-vector of the light. We can see from the above equation that, as $\Delta \rightarrow 0$, F_{scatt} increases. As the atoms will have some temperature, and therefore some velocity v , we can include the Doppler shift of the atoms (Eq. 5.1) into the detuning to give $\Delta = \omega_l - \omega_o + kv$, where ω_l is the frequency of the laser. Therefore, if we detuned the light slightly below resonance (i.e., red-detuned), then atoms travelling in the opposite

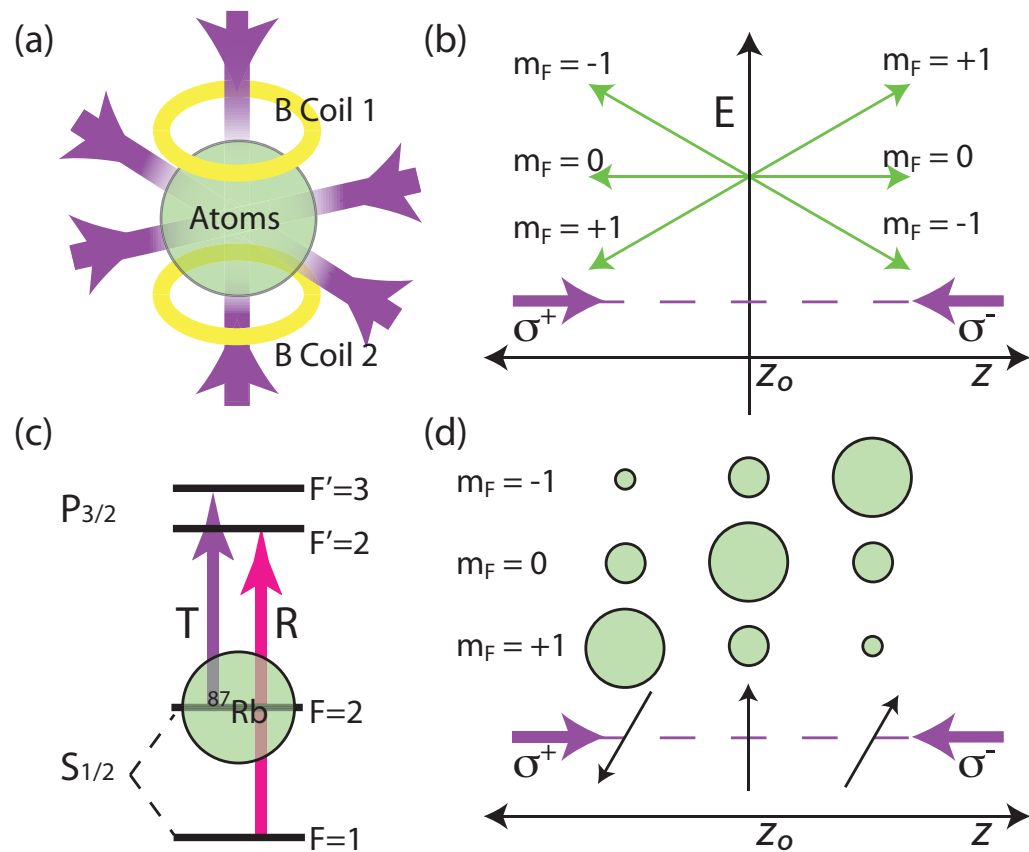


Figure 9.1: Magneto-Optical Trap. (a) The basic set-up for a MOT: purple arrows indicate the direction of propagation for the optical trapping fields; B Coil 1 and 2 show the position for the anti-Helmholtz coil configuration; and Atoms indicates where the atoms are trapped. (b) The principle behind the operation of a MOT: E is the energy of the internal atomic m_F states (green) and trapping lasers (purple) as a function of position z from the centre of the trap z_0 . Polarisation of trapping fields also shown. Note: zero energy is not shown. (c) Level scheme for trapping (T) and repump fields (R) in a ^{87}Rb MOT, with the green circle showing where the atomic population resides. (d) Polarisation-gradient cooling scheme: purple arrows show direction of propagation of the trapping fields; black arrows indicate the direction of ellipticity for the combined optical field (and therefore quantisation axis); circles indicate the relative atomic populations for the three m_F states at different positions along the z -axis from the origin at z_0 as labelled (not to scale).

direction to the light field will be Doppler shifted towards resonance and experience a greater force against their motion.

If we combined three sets of orthogonal counter-propagating light fields, atoms will be slowed whenever they move away from the centre of the fields. This is the basis of molasses cooling, first demonstrated by S. Chu et al. in 1985 [354], with a force in any direction given by [204]

$$F_{mol} \approx -2 \frac{\partial F_{scatt}}{\partial \omega} kv. \quad (9.2)$$

There are a couple of issues with molasses cooling, the first being that it is not a trap - as there is no restoring force towards the centre - and atoms will move out of the cooled cloud after a few seconds. Also the capture range, i.e., the range of atomic velocities that can be cooled, is too small to cool an initially thermal cloud of atoms. Finally, as there will be heating due to the random direction of spontaneous emission, there will be a minimum temperature for atoms cooled in a optical molasses.

To address the first issue, and create a restoring force, a magnetic field with a minimum at the centre of the ensemble can be added to the optical molasses. This can be achieved using a pair of coils in an anti-Helmholtz configuration. Combining this with using correct σ^\pm polarisations for the counter-propagating light fields will lead to trapping of the atoms. This MOT set-up is shown in Fig. 9.1(a). The restoring force is created by preferential absorption of one of the trapping fields over the other: the addition of the magnetic field gradient leads to Zeeman splitting of the m_F states away from $z = 0$, bringing one m_F state closer to resonance with the red-detuned optical fields on either side of the centre of the trap. For example, let us assume that in the positive z -direction the $m_F < 0$ states are at lower energy. Then, if we choose the polarisation of the light field propagating in the negative (positive) z -direction to have σ^- (σ^+) polarisation, the atoms in the $m_F < 0$ (> 0) states will preferentially absorb light from the σ^- (σ^+) field (due to Clebsch-Gordon coefficients - see Ref. [184]) and therefore receive a force towards the centre of the trap. This force will become larger the further the atom moves away from the centre as it will bring the atom closer to resonance with the trapping fields. This is illustrated in Fig. 9.1(b). The strength of the force on the atoms in a MOT is given by [204]

$$F_{MOT} = F_{mol} - 2 \frac{\partial F_{scatt}}{\partial \omega} \frac{g_F \mu_B}{\hbar} \frac{dB}{dz} z, \quad (9.3)$$

where g_F is the Landé g -factor, μ_B is the Bohr magneton and dB/dz is the magnetic field gradient in the z -direction. It is worth noting that, due to interatomic collisions, many atoms that are initially in a $m_F \geq 0$ state and stray in the positive z -direction will at some stage end up in a $m_F < 0$ state, and therefore feel a restoring force, while always undergoing molasses cooling.

In reality there is one more effect that must be dealt with for either of the above techniques to function: optical pumping. With an alkali element there are always two ground levels. Taking ^{87}Rb as an example, if we were to use the $F = 2$ level for the trapping state, and therefore have a laser slightly detuned from a $F = 2 \rightarrow F'$ transition, then the atoms have some probability of decaying into the $F = 1$ level, where they would be invisible to the trapping fields and can therefore escape the trap. Even if we used the closed D2 $F = 2 \rightarrow F' = 3$ transition atoms would still end up being lost too quickly to allow a trap to form. The solution to this is to add another field, a repump, to push atoms back from the $F = 1$ to the $F = 2$ ground level. This is illustrated in Fig. 9.1(c).

The first experimental demonstration of a MOT was by Raab et al. in 1987 [355] for a sample of cooled Na atoms. Thermal atoms can also be loaded directly into a MOT as the restoring force allows for a much larger capture range compared to an optical molasses. In practise, a MOT is usually used to trap thermal atoms before the magnetic field is turned off to allow for more optical-only cooling, as discussed in the following section. It is also worth noting that the same optical and magnetic fields that allow us to trap atoms mean that we can't store light using a MOT. The optical fields are detuned by only a few γ below resonance and the scattering rates associated with these (see Sec. 9.3.3), as well as having magnetic field gradients on the order of 10 G cm^{-1} [356] will lead to decoherence rates far too large for a useful quantum memory. Instead, a normal experimental protocol will involve turning off both magnetic and optical fields before any storage is attempted. Other, conservative, traps can be used to contain the atoms and these will be discussed in Sec. 9.2.3.

9.2.2 Further Cooling

Turning off the magnetic fields after a MOT has been loaded can lead to further cooling of the atoms. This is achieved via polarisation-gradient cooling (PGC). The configuration for PGC is the same as that for molasses cooling, i.e., three sets of orthogonal, counter-propagating beams, but in this case the beams are adiabatically detuned further from resonance.

There are two different polarisation combinations for PGC, with the beams either being orthogonal-linear or orthogonal-circular. The outcome in both cases is the same. However, the mechanics of the two cases are very different. As we will be using the MOT trapping fields with orthogonal-circular polarisations, we will discuss only this case here. Both cases are presented in detail in Ref. [357].

If we consider the beating of the σ^+/σ^- fields, the positive components of the electric fields for the separate polarisations, and total electric field, are given by

$$\begin{aligned}
 E^+(z) &= -\frac{\mathcal{E}^+}{\sqrt{2}} (\epsilon_H + i\epsilon_V) e^{ikz} \\
 E^-(z) &= \frac{\mathcal{E}^-}{\sqrt{2}} (\epsilon_H - i\epsilon_V) e^{-ikz} \\
 E_t(z) &= \frac{1}{\sqrt{2}} (\mathcal{E}^- - \mathcal{E}^+) (\cos(kz)\epsilon_H - \sin(kz)\epsilon_V) \\
 &\quad - \frac{i}{\sqrt{2}} (\mathcal{E}^- + \mathcal{E}^+) (\sin(kz)\epsilon_H + \cos(kz)\epsilon_V), \tag{9.4}
 \end{aligned}$$

where \mathcal{E}^+ (\mathcal{E}^-) are the amplitudes for the σ^+ (σ^-) polarised fields, and ϵ_H (ϵ_V) are the unit polarisation vectors for horizontal (vertical) light travelling in the positive (negative) z -direction. As can be seen from Eq. 9.4, the combined electric field will be elliptically polarised, with ellipticity $(\mathcal{E}^- - \mathcal{E}^+) / (\mathcal{E}^+ + \mathcal{E}^-)$. The direction of ellipticity will change angle according to $\phi = -kz$. If $\mathcal{E}^+ = \mathcal{E}^-$, as is normally the case, then the light will always be linearly polarised and vary between all-horizontal and all-vertical polarised light in a helical pattern.

The polarisation of the light will define a quantisation axis for an atom in this combined

field. If, for ease of explanation, we assume that the transition of interest is a ^{87}Rb D2 $F = 1 \rightarrow F' = 2$ transition, then it follows that the optical pumping rate for atoms from $m_F = \pm 1 \rightarrow m_F = 0$ will be $1/4$ while the pumping rate from $m_F = 0 \rightarrow m_F = \pm 1$ will be $1/9$ for the π -polarised light (again due to Clebsch-Gordon coefficients - see Ref. [184]). This will lead to a steady-state population for an ensemble with the majority of atoms in the $m_F = 0$ state and equal numbers in the two $m_F = \pm 1$ states. If the atom moves along the beams, the direction of polarisation and therefore the quantisation axis will change. It can be shown that this will not change the level structure, but will alter the relative number of atoms in the $m_F = \pm 1$ states. This difference in populations for an atom moving at speed v is given by [357]

$$\langle m_F = +1 \rangle - \langle m_F = -1 \rangle = \frac{40}{17} kv\Delta, \quad (9.5)$$

where Δ is again the detuning of the light from resonance. As long as Δ is negative, atoms moving in the positive z -direction will have a greater population in the $m_F = -1$ state. Now, we know from above that atoms in $m_F = -1$ are more likely to absorb a σ^- photon than a σ^+ photon and so, if the directions and detuning are correct, this will lead to another friction-like force applied to the atoms. This process is illustrated in Fig. 9.1(d). The temperature of the PGC sample with potential energy U_o will be given by [204]

$$T_{pgc} \approx U_o/k_B \propto \frac{I}{|\Delta|}, \quad (9.6)$$

with the minimum temperature of the ensemble now determined by the recoil limit

$$T_{rec} = \frac{\hbar^2 k^2}{k_B M}, \quad (9.7)$$

where M is the mass of the atom and k_B is the Boltzmann constant. Minimum temperatures achieved with PGC can be down to microkelvin for Cs in a 1D optical lattice [358].

The temperature of atoms can also be cooled past the recoil limit, down to tens of nanokelvin. This is achieved by evaporative cooling: where the hotter atoms are removed from the ensemble, followed by a rethermalisation of the remaining atoms [359]. This rethermalisation is a competition between “good” (i.e., elastic) energy transferring collision and “bad” (i.e., inelastic) heating collisions. There is no fundamental limit for evaporative cooling, though more evaporation will result in less atoms in the final sample. Other sub-recoil-limit cooling techniques involve velocity-selective coherent population trapping (VSCPT) [360] or Raman cooling [361]. Cooling atoms down to these temperatures can result in the creating of a Bose-Einstein condensate (BEC) [362, 363, 364], where the atoms become degenerate (i.e., all in the same quantum state). The incredibly high densities, low temperatures and quantum properties of a BEC mean that it is also a possible medium for a quantum memory. However this must be balanced with the fact that there are orders of magnitude less atoms in a BEC than in a MOT.

9.2.3 Dipole Traps and Optical Lattices

An alternative to simply turning off the trapping fields of a MOT for quantum memory operation is to shine a high intensity, far-detuned laser onto the MOT [365]. As the laser




Figure 9.2: Dipole Trap Modelling. Dipole trap depth U_d (blue solid line, left-hand scale) and dipole force F_d (red dashed line, right-hand scale normalised to the force of gravity F_g) (a) as a function of position from the centre of the trap beam r with waist w_d ; (b) as a function of position from the focus z for a single focused beam; and (c) as a function of z for a standing-wave dipole trap. For simulations, 1 W of power was assumed, with $w_d = 10 \mu\text{m}$ and $\lambda_d = 1064 \text{ nm}$ (i.e., $\Delta_d = -1 \cdot 10^{14} \text{ Hz}$ for the $F = 1$ level of ^{87}Rb).

is far-detuned there will be very few absorption-reabsorption events, making it much more promising for quantum memory purposes. Instead the interaction will mostly be reactive, with the laser causing a shift in the energy levels of the atoms. This is known as the ac Stark effect [366, 367], which will be discussed in more detail in Ch. 10. The change in energy of the internal states will create either a potential well (red-detuned light) or hill (blue-detuned light) and, therefore, the laser will either attract or repel the atoms, respectively, by the dipole force. Thermal atoms cannot be loaded directly into this type of trap as it is conservative and, therefore, can only be used to trap atoms with energies below the height of the trapping potential barrier. The size and direction of the dipole force, which determines the strength and the depth of such a far-off-resonance trap (FORT), depends on the gradient of the potential U_d caused by the optical field following [204]

$$\mathbf{F}_d(\Delta, r, z) = -\nabla U_d(\Delta, I(r, z)) = -\nabla(\bar{U}_d(\Delta) \cdot I(r, z)) \approx -\bar{U}_d(\Delta) \partial_r I(r) \mathbf{r}, \quad (9.8)$$

with $I(r, z)$ being the intensity of the dipole field travelling in the z -direction. The dependence of U on Δ will be discussed in Ch. 10, and \bar{U} denotes the potential for unit intensity. The approximation assumes a Gaussian beam profile, as the main variation in the beam will be radially outwards ($r^2 = x^2 + y^2$), rather than along the z -axis, with \mathbf{r} being the radial unit vector. A consequence of this is that the trapping in the z -direction will be much weaker than in the radial trapping.

Figures 9.2(a) and (b) show U_d and F_d across (r) and along (z) a single-beam dipole trap with a focussed waist of $w_d = 10 \mu\text{m}$ and a wavelength of $\lambda_d = 1064 \text{ nm}$. The restoring nature of the force in the r -direction can be seen from Fig. 9.2(a), with the maximum force occurring at $r = w_d/2$. In comparison, Fig. 9.2(b) shows the restoring force in the z -direction for the same trap. Though the trapping force is a factor of 50 smaller in the z -direction, a single focused beam was used by S. Chu et al. [368] for the first demonstration of a red-detuned dipole trap, with approximately 500 Na atoms trapped for several seconds. For longer lifetimes one could consider using a hollow-core blue-detuned dipole trap [369, 370]. Here the atoms will be repelled from the region of high intensity and, therefore, the overall scattering rate will be lower. However, here there will be even less confinement in the z -direction.

One option for improving the trapping in the z -direction is to use multiple trapping beams. A standing-wave dipole trap, where two dipole beams are sent counter-propagating, will lead to trapping nodes spaced by $\lambda_d/2$ [371]. This is shown in Fig. 9.2(c), where the total laser power is the same as that used for Fig. 9.2(b). As can be seen, due to the fast oscillations in the ground state energy, the relative trapping strength in the z -direction is much greater for the standing-wave configuration. However, the tight confinement in the z -direction will lead to oscillation of the atoms and heating effects which will be discussed in Sec. 9.3.4. Another option is to use a crossed-beam dipole trap [372], where two focused trapping beams are crossed at an angle. The greater the angle between the two beams, the less axial confinement can be achieved.

The loading efficiency from a MOT to a dipole trap can be up to 100% for individual atoms [373], but normally the number of atoms loaded into a dipole trap cannot exceed 40% of the atoms in a MOT [374], leading to atom numbers on the order of 10^7 at $100 \mu\text{K}$ [375]. Dipole traps have lifetimes on the order of hundreds of milliseconds [376] to seconds [377], normally limited by collisions and heating.

By adding more dipole beams, either red- or blue-detuned, optical lattices can be created. However, these are outside the scope of this thesis. For more information on

these see, for example, Ref. [378]. With optical lattices there will be a balancing act between increased confinement and reduction in atom number due to less trapping regions.

9.3 Decoherence Effects in Cold Atomic Ensembles

In this section we present the main decoherence mechanisms present in cold atomic ensembles.

9.3.1 Collisions

The collision rate in cold atomic ensembles has been studied in depth (see, for example, Ref.s [379, 380, 381]) with the collisional loss rate for a trap with an atomic density n being given by

$$\frac{dn}{dt} = -\alpha n - \beta n^2. \quad (9.9)$$

As can be seen from the above equation, there are two components to the decay: one due to background gas collisions with coefficient α being determined by the background gas pressure; and one due to inelastic collisions between atoms in the trap with coefficient β . Estimates of $1/\alpha$ vary from approximately one second for a dipole trap for a trap pressure of $3 \cdot 10^{-9}$ mbar [378], up to 20 s for Rb [379] and 132 s for Cs [380].

The β term is perhaps more interesting as it can give us an idea of the timescale for inelastic collisions between atoms in the trap. This is slightly complicated by the fact that, in general, β is not the total rate of collisions, but the rate of collisions that can cause atoms to leave the trap and is therefore dependent on the trap depth (as the shallower a trap, the lower the energy of the collisions needed for the atoms to leave). Wallace et al. investigated β as a function of dipole laser intensity (which determines trap depth $U_d \propto I$) for ^{87}Rb and found that $\beta \rightarrow 2 \cdot 10^{11}$ cm³/s as $I \rightarrow 0$ [379]. This is the total hyperfine-changing collision rate, a lower-bound on the total collision rate (as there are also non-hyperfine-changing collisions to take into account for memory decoherence purposes).

9.3.2 Diffusion and Time of Flight

If atoms are released from their trap, the probability distribution of the ensemble after time t (assuming x is the vertical axis and therefore affected by gravity with constant $g = 9.8$ m s⁻²) is given by [368]

$$N(x, y, z, t) = \text{Gs} \left[x - \frac{gt^2}{2}, \sigma_t \right] \text{Gs}[y, \sigma_t] \text{Gs}[z, \sigma_t], \quad (9.10)$$

where $\text{Gs}[x, \sigma]$ is a Gaussian function of the form $\text{Gs}[x, \sigma] = (2\pi\sigma^2)^{-1/2} \text{Exp}[-x^2/2\sigma^2]$ and $\sigma_t = \sqrt{\sigma_o^2 + \sigma_v^2 t^2}$ is the Gaussian radius of the ballistically expanded cloud, with σ_o the initial spherical radius of the cloud (i.e., used to determine trap velocity). This all depends of the temperature of the cloud by

$$T = \frac{M}{k_B} \sigma_v^2. \quad (9.11)$$




Figure 9.3: Cold Atom Diffusion. The fraction of atoms remaining in interaction region in (i) the vertical (x) direction (blue); (ii) the horizontal (y) direction (red); (iii) the axial (z) direction (green); and (iv) total (orange). The shape of the ensemble was taken to be cigar shape, with $\sigma_{x,y} = 1/6$ mm and $\sigma_z = 1/6$ cm, with $T = 300$ μ K.

If we define the interaction region to be a cylinder (for reasons that will become apparent later¹) of radius $R = 1$ mm in the x, y -plane and length $L = 1$ cm in the z -direction we can take $\sigma_{x_o} = \sigma_{y_o} = R/6$ and $\sigma_{z_o} = L/6$, to ensure 99% of the atoms lie within six standard deviations of the mean. If we define this as the interaction area, we can model the diffusion over time. Figure 9.3 shows the fraction of atoms $N(t)/N(0)$ remaining in the interaction region over time, taking $T = 300$ μ K and assuming an initial constant density. As can be seen, diffusion rates in the radial-direction are much higher than in the z -direction. This is not surprising considering the much smaller size of the interaction area along these axes. The decoherence above is Gaussian, with decay of the form $\text{Exp}[-t^2/t_{diff}^2]$, where t_{diff} is the Gaussian decay constant due to diffusion. For the cloud modelled above $t_{diff} \approx 2.5$ ms.

We have already discussed the effect of longitudinal diffusion in GEM (see Sec. 5.3.2). One way of overcoming this issue with cold atoms would be to use a standing-wave dipole trap or optical lattice, as described in Sec. 9.2.3. However, these methods will lead to increased heating rates, discussed in Sec. 9.3.4 below.

9.3.3 Scattering Rate

Another source of decoherence, as well as trap loss, is the scattering rate of the atoms due to absorption and re-emission of light. The full equation for the scattering rate Γ_{sc} from an initial state $|g_i\rangle = |F_i, mF_i\rangle$ with laser frequency ω_l is [207, 382]

$$\Gamma_{F,m_F} = \frac{I}{6\pi\epsilon_0^2\hbar^3c^4} \sum_{g_f} (\omega_l - \omega_{f_i})^3 \times \left| \sum_{a,q_{sc}} \frac{\langle g_f | \mathbf{e}\mathbf{r} \cdot \boldsymbol{\epsilon}_{q_{sc}} | a \rangle \langle a | \mathbf{e}\mathbf{r} \cdot \boldsymbol{\epsilon}_q | g_i \rangle}{\omega_{ag_i} - \omega_l} \right|^2, \quad (9.12)$$

¹For instance, see Sec. 9.4.2 or Ch. 12.




Figure 9.4: Scattering Rate. Scattering rate ($/2\pi$) for the $F = 2$, $m_F = 0$ state of ^{87}Rb with unit intensity for: (a) small detunings ($\Delta \approx \gamma$); (b) medium detunings ($\Delta \approx \Delta_{hfs}$); and (c) large detunings ($\Delta \gg \Delta_{hfs}$) with a circularly polarised light field.

in the rotating-wave approximation, where $|g_f\rangle$ is the final state (which cannot be higher in energy than $E_{g_i} + \hbar\omega_l$ by conservation of energy), ω_{fi} is the frequency of the transition from state $|g_i\rangle$ to $|g_f\rangle$ (negative if $|g_f\rangle$ is lower than $|g_i\rangle$ in energy), $|a\rangle$ is the dummy state representing all possible intermediate states (not restricted by energy constraints), q is the polarisation ($-1, 0, +1$) of the initial photon, q_{sc} is the polarisation of the scattered photon, and ω_{agi} is the frequency of the transition between states $|g_i\rangle$ and $|a\rangle$. The unit-vectors $\epsilon_{q/q_{sc}}$ and \mathbf{r} are the direction of the electric field due to the polarisation q/q_{sc} and the direction of the electric dipole moment of the atom, respectively. The matrix elements can be calculated (see App. D) using the above equation, or it can be simplified if the detuning is much greater than the excited state hyperfine-splitting Δ'_{hfs} , to give

$$\Gamma_{F,m_F} = \frac{I}{6\pi\epsilon_0^2\hbar^3c^4} \sum_{g_f} (\omega_l - \omega_{fi})^3 \left| \frac{A_{1/2,g_i}}{\Delta_{1/2}} + \frac{A_{3/2,g_i}}{\Delta_{3/2}} \right|^2 = \bar{\Gamma}_{F,m_F}(\Delta) I(r, z), \quad (9.13)$$

where $\Delta_{J'}$ is the detuning of the laser from the J' transition, $\bar{\Gamma}_{F,m_F}$ is the scattering rate per unit intensity, and

$$A_{J',g_i} \equiv \sum_{a,q_{sc}} \langle g_f | \mathbf{e}\mathbf{r} \cdot \epsilon_{q_{sc}} | a \rangle \langle a | \mathbf{e}\mathbf{r} \cdot \epsilon_q | g_i \rangle \quad (9.14)$$

for all states $|a\rangle$ within the level J' level. Figure 9.4 shows scattering rates calculated with Eq. 9.12 for the $F = 2$, $m_F = 0$ state for circularly polarised light for detunings close to resonance ($\Delta \approx \gamma$, as is the case in a MOT), as well as around the ground-state

hyperfine-splitting ($\Delta \approx \Delta_{hfs} = 6.8$ GHz, see App. A) and far-detuned ($\Delta \approx \Delta_{hfs}$, as is the case in a dipole trap).

It should be noted that there are two different types of scattering: Raman (or inelastic) scattering, where the final state is different from the initial state, and Rayleigh (or elastic) scattering, where the initial and final states are the same. From Eq. 9.14 it can be found that for Raman scattering $A_{3/2} \approx -A_{1/2}$ and for Rayleigh scattering $A_{3/2} \approx 2 \cdot A_{1/2}$. When $\Delta \gg \Delta_{fs}$ and therefore $\Delta_{1/2} \approx \Delta_{3/2}$ it can be seen from Eq. 9.13 that scattering between levels will be suppressed. From experiments in dipole traps it has been shown that the inelastic scattering rate is slowed by an additional factor of $1/\Delta^2$, extending the $F = 2 \rightarrow F = 1$ relaxation rates by orders of magnitude [377]. Unfortunately this is unlikely to decrease the decoherence rate, as Rayleigh scattering will also cause loss of coherence.

9.3.4 Heating Rate

One factor that may play a big part in the lifetime of a dipole trap (though not necessarily the coherence time) is the rate at which the atoms heat \dot{T} , as once they have reached a certain energy, $E = E_{pot} + E_{kin} > U_d$, the atoms will be able to escape. The heating rate can be given in terms of the recoil temperature T_{rec} (i.e., the temperature gained by the atom, in terms of kinetic energy, by emitting one photon), κ which is defined as the ratio of potential to kinetic energy (equals one for a 3D harmonic trap and zero for a box potential), and the scattering rate determined above to be [378]

$$\dot{T} = \frac{2/3}{1 + \kappa} T_{rec} \Gamma'_{sc}, \quad (9.15)$$

where Γ'_{sc} is the mean scattering rate, with T_{rec} defined in Eq. 9.7. For a dipole trap at 1064 nm with 20 W of power focused to a 150 μm waist, the heating rate will be approximately 55 nK s^{-1} , with a trap depth of 75 μK .

Dipole traps will also heat if there are intensity fluctuations or beam jitter. In a standing-wave dipole trap, or optical lattice, there will also be heating due to phase jitter between two counter-propagating fields, as this will shake the atoms. For more information on these effects see Ref.s [383, 384].

9.4 Literature Review of Cold Atom Light Storage Experiments

In this section we will review some of the progress made so far towards developing a cold atom quantum memory. This includes both a history of cold atom quantum memory experiments, as well as a description of some of the techniques used to create the coldest and densest ensembles possible.

9.4.1 A History of Cold Atom Light Storage Experiments 1999-Present

The idea of using cold atoms as a storage medium for quantum information protocols has been around for nearly as long as the protocols themselves. For instance, one of the first

demonstrations of EIT, where light was slowed down to 17 ms^{-1} , was achieved in a Na BEC by the group of L. V. Hau et al. in 1999 [118]. This group was able to show a pulse delay of approximately $7 \mu\text{s}$ with $1\text{-}2 \cdot 10^6$ atoms in their BEC, with an EIT feature full-width-half-maximum (FWHM) of 1-2 MHz. Following from this work, in 2001 the same group demonstrated classical light storage using EIT, in a sample of atoms cooled to just above the critical temperature for a BEC, with a coherence time (i.e., $1/e$ decay time) of approximately 1 ms and a storage efficiency of 1-2% [121] for a storage time of just over 1 ms.

In 2008 more progress was made, with the group of B. Zhao et al. demonstrating quantum storage using cold ^{87}Rb with a coherence time over 1 ms [385]. They achieved this by storing the information using a DLCZ-style scheme in a sample of approximately 10^8 atoms, with an optical depth (OD) of five, released from a MOT. The limit of storage in this case was the diffusion of the now untrapped atoms from the interaction region.

In the same year, and indeed in adjacent papers, the group of R. Zhao et al. also demonstrated quantum storage in a cold atomic sample [386]. In this case it was with $10^5\text{-}10^6$ ^{87}Rb atoms trapped in a 1D optical lattice, with the information again being stored via a DLCZ-style scheme. The optical lattice prevented the diffusion of the atoms, with a coherence time of up to 6 ms observed, though with efficiencies of only a few percent.

In 2009, the group of Y. O. Dudin et al. demonstrated storage of entanglement using 10^7 ^{87}Rb atoms trapped in a 2D optical lattice [387]. Here the group was able to demonstrate a violation of Bell's inequality after storing a photon for 3 ms, using a DLCZ-style Raman scattering technique. The following year, this group was also able to demonstrate coherent light storage with a coherence time of up to 320 ms [388]. This was achieved with a 1D optical lattice, again made up of 10^7 ^{87}Rb atoms, with efficiencies of a few percent.

Back in 2009, two other papers both published results showing long storage times using cold atoms. The first paper was from the group of U. Schnorrberger et al., who stored classical light with a coherence time of 240 ms [389]. This was achieved using EIT with a sample of approximately 10^5 ^{87}Rb atoms trapped in a Mott insulator. This sample gave an optical depth of around six, with efficiencies less than one percent. The EIT feature in this case had a FWHM of approximately 80 Hz. The second paper was, again, from the group of L. V. Hau. This time, using a Na BEC stored in a dipole trap with approximately $3 \cdot 10^6$ atoms, a coherence time of 540 ms was achieved using EIT [119], though again the efficiencies were on the order of one percent. As can be seen, these results are starting to approach the coherence times of 2.3 s achieved with solid state ensembles [5].

High efficiencies have also been achieved in cold atomic ensembles. For example, S. Zhang et al. demonstrated up to 50% efficiency for EIT with a 200 ns storage time in a 2D ^{85}Rb MOT with an OD greater than 50 [390]. The efficiency was limited to 50% by four-wave mixing processes. Another group also demonstrated approximately 50% efficiency for EIT, this time for single-photons in a ^{87}Rb BEC in a cross-beam dipole trap containing $1.2 \cdot 10^6$ atoms [391]. The coherence time for this experiment was approximately $500 \mu\text{s}$. The 50% efficiency limit was recently beaten by the group of Y.-H. Chen et al., with 78% efficiency and a coherence time of $98 \mu\text{s}$ using a MOT with 10^9 ^{87}Rb atoms and an OD of 160 [13]. With this system they were able to demonstrate a delay-bandwidth product of 74 over 50% efficiency.

It has also been demonstrated by the group of J.-W. Pan that, by creating a MOT inside a ring cavity, single-photons can be generated efficiently using a DLCZ scheme, with the atomic coherence being maintained for up to 3 ms [149]. Finally, experiments have also been done showing the interaction of quantum states of light with cold-atom-based

EIT. Continuous-wave sideband squeezing [131] and discrete-variable entanglement [103] have both been investigated.

9.4.2 Experimental Improvements

Improving Lifetime

To try to improve storage times, some groups have come up with ingenious methods to overcome some of the decoherence effects discussed in Sec. 9.3. For instance, the group of Y. Sagi et al. in 2010 were able to decrease the decoherence effects due to interatomic collisions by applying more than 200 π -pulses to their 1D lattice of approximately $3 \cdot 10^5$ ^{87}Rb atoms [392]. The densities they were able to achieve gave them an OD of 230, and a 20-fold decrease in collisional decoherence gave them a coherence time of greater than 3 s, although no memory was demonstrated in this experiment.

Cooling the sample as much as possible, such as with polarisation-gradient cooling (Sec. 9.2.2), will increase coherence times for experiments where cold atoms are released from a MOT. However, in the end gravity will accelerate even the coldest cloud downwards. One elegant solution to this is to use the vertical direction for interrogation of the atoms, as was done by the group of J.-W. Pan in their experiment mentioned in the previous section [149].

Some groups, such as the group of F. Yang et al.² [393] and the group of A. G. Radnaev et al. [394] among others, have made use of compensation light fields to increase their storage lifetimes. This is based on the idea of A. Kaplan et al. that additional light fields can be used to decrease the broadening of spectroscopic lines caused by, for instance, dipole traps [395] or non-homogeneous background magnetic fields. In the case of F. Yang, they were able to see a 40-fold increase in coherence time for their hollow dipole trap, giving them an EIT coherence time of nearly 30 ms. Radnaev et al. were able to produce efficiencies of up to 7% using a compensation beam for their 1D lattice and a modified DLCZ-style wavelength-conversion memory, with coherence times of 165 ms achieved.

Improving Optical Depth

The brute force method for improving the number of atoms in a MOT is to make the trapping beams as large as possible to maximise the trapping volume, with as much power as possible. However, in the end it is the optical depth that will determine the efficiency of a cold atomic memory, and this depends on the density of atoms in the interaction region.

One optimisation issue with a MOT is the background vapour pressure. To have high atom numbers in a standard MOT normally requires quite a dense background vapour (normally on the order of 10^{-9} Torr for Rb). However, in Sec. 9.3.1 we discussed how collisions with background atoms can reduce the atom number in a MOT as α increases. This will lead to an optimal value for the background pressure: too low and the MOT will not have a large enough background reservoir to fill from, especially if the load time is relatively short; and too high and the collisions will limit the size of the MOT. One way around this is to have a two-chamber MOT, with high background pressure in the first chamber and a low pressure in the second chamber, achieved using an impedance, such as a narrow tube, to allow for a pressure differential. Atoms can then be transported from the first to the second chamber by either optical or magnetic fields. This was first

²The same group as B. Zhao et al. mentioned previously.

demonstrated by C. J. Myatt et al. in 1996 when they achieved a 90% transfer efficiency between their two chambers to obtain a ^{85}Rb MOT with $1.5 \cdot 10^{10}$ atoms in the second chamber [396]. This was $30\times$ more atoms than they held in the first chamber.

In 2008 the group of Y.-W. Lin et al. came up with a method for increasing the OD of a MOT without increasing the number of atoms present [397]. They achieved this by using rectangular, rather than circular, anti-Helmholtz coils. This creates a cigar-shaped MOT, placing more atoms in the interaction region. The dimensions of this MOT were initially $1 \times 2 \times 7.5$ mm with approximately $7 \cdot 10^7$ ^{87}Rb atoms present and an OD of 32. This was improved upon with further work in 2009 to give dimensions of $2.5 \times 2.5 \times 10$ mm with 10^9 atoms and an OD of 36 [398]. EIT was attempted with both set-ups, with an 8-10 μs pulse delay seen for the earlier work.

Another way of improving the OD without increasing atom number, as well as cooling the MOT further, is to perform a dynamic magnetic field compression before turning off the MOT. The reason this is performed dynamically is because trapping with a higher magnetic field gradient will lead to less atoms trapped and, therefore, no improvement in density. For instance, by optimising the initial magnetic field gradient to 11 G cm^{-1} and then ramping this up to 60 G cm^{-1} in 5 ms, accompanied by an increase in detuning of the trapping fields from -9 MHz to -44 MHz, the group of W. Petrich et al. were able to obtain a factor of 25 improvement in atomic density for a ^{85}Rb MOT, with a maximum density of $5 \cdot 10^{11} \text{ cm}^{-3}$ achieved [399].

Apart from polarisation-gradient cooling and compression, another way of increasing the density of the atoms in a MOT is to place a spatial dark spot on the repump field. This works because the density of the atoms in a MOT will reach a critical value in the trapping level (usually the higher ground level, for instance the $F = 2$ level for ^{87}Rb), due to the fact that light emitted by one atom will be absorbed by another while still within the MOT. This density limit is approximately 10^{11} cm^{-3} [356]. By placing a hole in the repump, atoms at the centre of the trap will fall into the untrapped level (i.e., the $F = 1$ level for ^{87}Rb) and not fluoresce or absorb light being emitted on the $F = 2$ level, allowing for a higher density of atoms in the centre of the trap. This was first demonstrated by W. Ketterle et al. in 1993 where they were able to achieve an increase in density from 10^{10} to 10^{12} cm^{-3} for their Na MOT [400]. The dark spot can also be temporal, as shown by S. Chu et al. in 1996 when they dramatically reduced their repump intensity 20 ms before turning off the trap [370]. This led to a temperature of 60 μK and a density of $6 \cdot 10^{11} \text{ cm}^{-3}$ for their Na MOT.

Combining the idea of a cigar-shaped MOT with a dark spot led the group of S. Zhang et al. to develop a high-OD atomic source with ^{85}Rb atoms in a 2D MOT and a dark line down the centre [401]. Though a 2D MOT, created with four trapping beams and a four-coil quadrupole magnetic field (for instance, see Ref. [402]), has an ideal geometry for quantum memory experiments, being very long in the z -direction and symmetric in the r -direction, by its nature it does not allow for cooling and trapping along the MOT. In this case, however, the authors used six trapping beams with the quadrupole coils to allow for cooling and trapping in all directions. Care was taken to ensure that four of the trapping beams were in the x,z -plane, with the final two beams travelling along the y -axis, and it was only along this axis that the repump beams were sent. This allowed for the placement of a dark line in the repump and, therefore, the creation of a region of high density, as discussed in the paragraph above, along the length of the non-symmetric ensemble. With this method, an OD of 160 was achieved using ^{85}Rb atoms, with a linear density of $2.05 \cdot 10^{15} \text{ m}^{-2}$.

In the end it is the density of atoms in the correct m_F ground state of the correct F level that will determine the efficiency for GEM. It is therefore important to ensure that as many atoms as possible are in this state. One way of achieving this is through optical pumping. Pumping from one F level to another is relatively easy - to place atoms in the trapping level the trapping field is turned off, and to place atoms in the non-trapped level the repump is turned off - but placing the atoms in the correct m_F state is more complicated. The group of B. Wang et al. were able to achieve this in a ^{87}Rb MOT by adding two additional fields either propagating with, or orthogonal to, the probe beam, and varying the polarisations of all three [403]. All fields were on the D1 line, with one on the $F = 1 \rightarrow F' = 1$ transition, another on the $F = 2 \rightarrow F' = 2$ transition and the probe on the $F = 1 \rightarrow F' = 2$ transition. With this set-up the authors were able to place atoms into any of the three $F = 1$ states with over 96% efficiency, measured using EIT lines.

9.5 Conclusions

In this chapter we have presented the basic theory necessary to understand the mechanics of cold atom trapping, as well as its limitations. We have also presented a comprehensive literature review showing the results already achieved in terms of storage of light using cold atoms and optimising cold atomic ensembles. These works show the promise of cold atom quantum memories - with long storage times and high efficiencies - which is the motivation behind our move towards cold atom GEM.

ac Stark Gradient Proposal

Carter: “This ion beam thingy, that is what, exactly?”

Henry: “Irradiated uranium isotopes...”

Carter: “Oh death ray. Why didn’t you say ‘death ray’?”

- “Dr. Nobel” *Eureka*, SyFy, 22 August 2006, TV Series.

10.1 Introduction - An Enlightened Idea

In Ch. 7 we discussed the utility of having a quantum memory that could perform coherent operations on stored information for applications such as nodes in quantum networks [48] and increased bit rates over quantum communication channels [293]. In that case we used a multi-element coil with warm atoms in a vapour cell to experimentally demonstrate a series of spectral manipulation operations. Even though we showed control over this gradient down to tens of kilohertz with that set-up, ultimately we would prefer to move away from using magnetic field (i.e., Zeeman) gradients for both Λ -GEM manipulation and field switching. The reason for this is that, compared to electric fields, magnetic fields are much harder to create and control. Transient magnetic fields, which occur during gradient switching due to the inductance of the coils, also limit the switching time and can affect the rephasing process.

Unfortunately a dc electric field (i.e., dc Stark) gradient is not an option for Λ -GEM, where we desire to split the hyperfine states by megahertz, as alkali elements (along with all free atoms) do not have a linear dc Stark shift [404] and, therefore, electric fields of hundreds of kilovolts would be needed [182]. The only option left to us then is an optical intensity (i.e., ac Stark - acS) gradient, which would allow for an all-optically controlled quantum memory. This gradient creation method was first put forward by Kraus et al. [177]. The move to cold atomic ensembles, proposed in the previous chapter, would prove beneficial for an acS gradient memory, as it is an intensity dependent effect and cold atoms can be persuaded to occupy small volumes, reducing the laser power necessary.

This chapter investigates the feasibility of using an acS-generated frequency gradient for Λ -GEM, in concert with a move to an ensemble of cold ^{87}Rb atoms. The rest of the chapter is structured as follows: the main body (Sec. 10.2) discusses the proposed experimental implementation, including the theory behind the ac Stark shift and how it would apply to the creation of a frequency gradient across an ensemble of alkali atoms, as well as the optimal experimental parameters. Then factors that may limit this scheme will be discussed in Sec. 10.3 such as coherence times and maximum efficiencies possible. Finally, in Sec. 10.4 we provide a small experimental demonstration of the creation of an acS gradient.




Figure 10.1: Proposed ac Stark GEM Experiment. (a) Envisioned set-up for Λ -GEM experiment using cold atoms and an ac Stark gradient with: acS - ac Stark gradient field (gold); P - probe field (red dashed); C - coupling field (red solid); DT - dipole trap field (orange); A - acousto-optic modulator; PC - Pockels cell; PBS - polarising beam-splitter; Q - quarter wave-plate; BSh - beam shaper; BS - 50:50 non-polarising beam-splitter; DC - dichroic mirrors (reflect at 1064 nm and transmit at 795 nm); ^{87}Rb - cold atomic ensemble. (b) Side-on view of the cylindrical atomic ensemble showing the direction of the intensity gradient $I(z)$ and radius R . (c)(i) Gaussian and (ii) linear intensity profiles for the acS field per unit power over the ensemble. Here L has been taken to be 1 cm and $R = 10 \mu\text{m}$, determined by the radius of the DT beam.

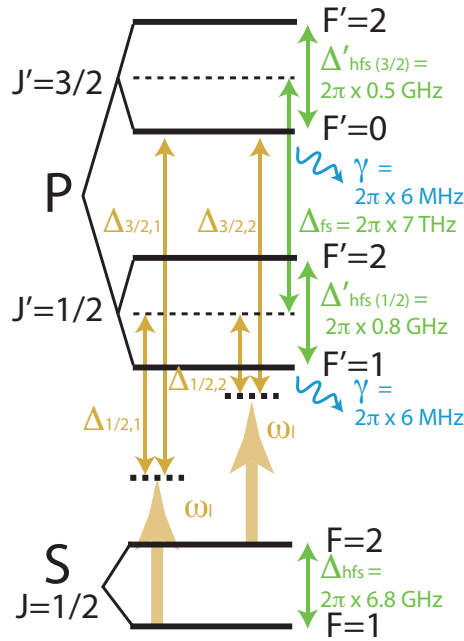


Figure 10.2: Rubidium-87 Level Structure. The level structure of rubidium-87 showing the hyperfine-splittings Δ_{hfs} between F levels for both ground and excited states, as well as the fine-structure-splitting Δ_{fs} between the two excited states (all green). Also shown are the detunings $\Delta_{J',F}$ for a given laser frequency ω_l (gold - see definitions in text), and excited-state decay rate γ (blue).

Most of the work presented in this chapter has been published in the following journal article:

- B. M. Sparkes, M. Hosseini, G. Hétet, P. K. Lam, and B. C. Buchler, *ac Stark gradient echo memory in cold atoms*, Physical Review A **82**, 043847 (2010).

10.2 AC Stark Shift Proposal

The envisaged experiment is shown in Fig. 10.1(a). The probe and coupling fields would be created similarly to the method described in Ch.s 6-8, with the probe field detuned an amount Δ from the D1 $F = 1 \rightarrow F' = 2$ transition and the coupling field 6.8 GHz away, Δ detuned from the D1 $F = 2 \rightarrow F' = 2$ transition.

There are three main components to this set-up: (i) cold atom storage; (ii) ac Stark gradient creation and (iii) switching (including probe and coupling field orientations). These will each be discussed in turn after firstly looking at the ac Stark theory.

10.2.1 AC Stark Shift Theory

When light of intensity $I(r, z)$, frequency ω_l and polarisation $q = 0, \pm 1$ - corresponding to linear, right (+) and left (-) circular polarisations respectively - is shone onto an atom there will be a change in energy of the internal states. This is known as the ac Stark (acS) effect [367, 366]. For an alkali atom, the basic structure of which is shown in Fig. 10.2 along

with definitions of the different detunings used below, this effect can be calculated for a given ground state $|g_i\rangle = |J = 1/2, F, m_F\rangle$ from second-order time-dependent perturbation theory to be [405]

$$U_{F,m_F}(\omega_l, q, I) = \frac{I(r, z)}{2c\epsilon_0\hbar} \sum_a \frac{|\langle a | e\hat{\mathbf{r}} \cdot \boldsymbol{\epsilon}_q | g_i \rangle|^2}{\omega_l - \omega_{ag_i}}, \quad (10.1)$$

where the sum is over all excited states $|a\rangle = |J', F', m'_F\rangle$, ω_{ag_i} is the frequency of the transition between states $|a\rangle$ and $|g_i\rangle$ allowing us to define the detuning as $\Delta = \omega_l - \omega_{ag_i}$. The unit-vectors $\boldsymbol{\epsilon}_{q/q_{sc}}$ and \mathbf{r} are the direction of the electric field due to the polarisation q/q_{sc} and the direction of the electric dipole moment of the atom, respectively. This formula uses the rotating-wave approximation, which is valid for detunings much smaller than the frequency of the transition, i.e., $\omega_l - \omega_{ag_i} \ll \omega_l + \omega_{ag_i}$. For simplicity we can write the above equation as $U_{F,m_F}(\Delta, q, I) = \bar{U}_{F,m_F}(\Delta, q)I(r, z)$, where \bar{U}_{F,m_F} is the change in energy per unit intensity. An approximation to Eq. 10.1, in the limit of $\Delta \gg \Delta'_{hfs}$, is given by [378]

$$U_{F,m_F}(\Delta, q, I) \simeq \frac{\pi c^2 \gamma I(r, z)}{2\omega_o^3} \times \left(\frac{2 + qg_F m_F}{\Delta_{3/2,F}} + \frac{1 - qg_F m_F}{\Delta_{1/2,F}} \right), \quad (10.2)$$

where γ (excited-state decay rate) and ω_o are the averaged values of the two excited levels, g_F is the Landé factor ($g_1 = -0.5$, $g_2 = -g_1$) and $\Delta_{J',F}$ is the detuning from the $S_{1/2}$ level ($F = 1, 2$) to either $J' = 1/2, 3/2$ levels, as illustrated in Fig. 10.2. This energy shift will remain linear in the limit $U \ll \hbar\Delta_{hfs}$.

Apart from changing the energies of the atomic levels, there is the possibility that the atoms will absorb and then re-emit the light. This scattering of light by the atoms will affect not only the coherence time τ_{coh} achievable with the ensemble, but also the lifetime τ_{trap} of any trap that is used to contain them (see Sec. 10.2.2). The scattering rate for a given ground state can be determined from Eq.s 9.12 and 9.13 in Ch. 9. As can be seen from these equations, $\bar{\Gamma}_{F,m_F} \propto 1/\Delta^2$ while $\bar{U}_{F,m_F} \propto 1/\Delta$ and therefore, in the context of laser trapping, increasing the detuning for a constant trap depth will increase both τ_{coh} and τ_{trap} .

Here we are interested in the energy splitting (and therefore frequency splitting $h\Delta\nu = \Delta U$) along an ensemble of atoms. This will determine the bandwidth of the system \mathcal{B}_s . The splitting per unit intensity between two adjacent m_F states in the same F level can be defined to be

$$\bar{\delta}_F(\Delta, |q|) \equiv 1/h |\bar{U}_{F,0}(\Delta, q) - \bar{U}_{F,1}(\Delta, q)|, \quad (10.3)$$

with the total magnitude of the splitting given by $\delta_F(\Delta, |q|, I) = \bar{\delta}_F(\Delta, |q|)I(r, z)$. Combining the above equation with Eq. 10.2 we obtain

$$\bar{\delta}_F(\Delta, |q|) = \frac{\pi c^2 \gamma}{2\omega_o^3 h} \left| \frac{qg_F}{\Delta_{1/2,F}} \right| \left(1 - \frac{\Delta_{1/2,F}}{\Delta_{3/2,F}} \right). \quad (10.4)$$

The frequency splitting between the $m_F = 0$ states of the $F = 1$ and 2 levels can similarly be found to be

$$\begin{aligned} \bar{\delta}_{1,2}(\Delta) &\equiv \frac{1}{h} (\bar{U}_{1,0}(\Delta) - \bar{U}_{2,0}(\Delta)) \\ &= \frac{\pi c^2 \gamma}{2\omega_o^3 h} \left(\frac{2(\Delta_{3/2,1} - \Delta_{3/2,2})}{\Delta_{3/2,2}\Delta_{3/2,1}} + \frac{\Delta_{1/2,1} - \Delta_{1/2,2}}{\Delta_{1/2,2}\Delta_{1/2,1}} \right) \end{aligned}$$

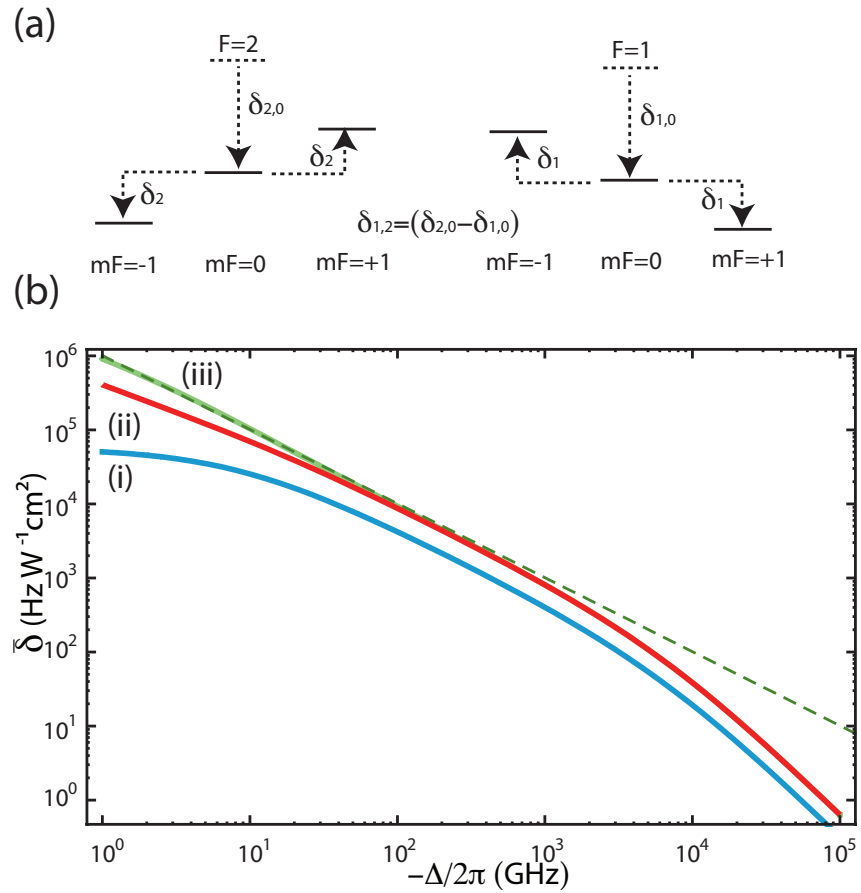


Figure 10.3: ac Stark Frequency Splitting. (a) Illustration of δ_F splitting for the $m_F = \pm 1, 0$ states of the $F = 1, 2$ levels. This is assuming a polarisation of $q = 1$ for the ac Stark field. $\delta_{1,2}$ is the difference between the splittings of the two $m_F = 0$ states of the respective levels from their original positions. (b) Log-log plot of (i) $\bar{\delta}_1$ (blue solid line), (ii) $\bar{\delta}_2 + \bar{\delta}_1$ (red solid line), and (iii) $\bar{\delta}_t = \bar{\delta}_2 + \bar{\delta}_1 + \bar{\delta}_{1,2}$ (green solid line) as a function of detuning $\Delta_{1/2,2}/2\pi$. The dashed green line shows a fit to $\bar{\delta}_t$ for small detunings where $\bar{\delta}_t \propto 1/\Delta$. These traces were calculated using Eq. 10.1.

$$= \frac{\pi c^2 \gamma \Delta_{hfs}}{2\omega_o^3 \hbar} \left(\frac{2}{\Delta_{3/2,2} \Delta_{3/2,1}} + \frac{1}{\Delta_{1/2,2} \Delta_{1/2,1}} \right). \quad (10.5)$$

These frequency splittings are illustrated in Fig. 10.3(a) with the polarisation taken to be $q = 1$ (i.e., right circularly polarised) as, if linear polarisation were to be used, $\bar{\delta}_F = 0$.

If, as was the case in the previous warm atom Λ -GEM experiments, one m_F state from each of the hyperfine levels is used for the ground states $|1\rangle$ and $|2\rangle$ of the Λ system, then we can define the total splitting in terms of $\bar{\delta}_F$ and $\bar{\delta}_{1,2}$ to be

$$\bar{\delta}_t(\Delta, q) \equiv \bar{\delta}_{1,2}(\Delta) - q (m_2 \bar{\delta}_2(\Delta, |q\rangle) + m_1 \bar{\delta}_1(\Delta, |q\rangle)). \quad (10.6)$$

Figure 10.3(b) shows the absolute value of δ_t for $m_2 = -m_1 = -1$ and the relative contributions from the three terms above as a function of detuning for $q = 1$. As can be seen, at large detunings $\Delta_{1/2,F} \gg \Delta_{hfs}$ the splitting between m_F levels becomes approximately equal (i.e., $\bar{\delta}_1 \sim \bar{\delta}_2$) and the relative contribution of $\bar{\delta}_{1,2} \rightarrow 0$. This behaviour can be explained from Eq.s 10.4 and 10.5 above, as $\Delta_{J',1} \simeq \Delta_{J',2}$ for $\Delta_{1/2,F} \gg \Delta_{hfs}$.

From Eq. 10.6 we can calculate the field gradient to be $\eta(z) \simeq 2\pi \bar{\delta}_t \partial_z I(r, z)$, assuming negligible change for the intensity in the transverse ($\hat{\mathbf{r}}$) direction. Using this equation for η allows the system bandwidth to be expressed in terms of $\bar{\delta}_t$ and I as follows

$$\begin{aligned} \mathcal{B}_s(\Delta, |q\rangle I) &= \frac{1}{2\pi} \int_0^L dz |\eta(\Delta, q, I)| \\ &= |\bar{\delta}_t(\Delta, q) (I(z=L) - I(z=0))|. \end{aligned} \quad (10.7)$$

10.2.2 Atom Trapping

As discussed in the previous chapter, a magneto-optical trap (MOT) is the most common method for cooling and trapping atoms. However, due to scattering of light caused by the near-resonant lasers involved and the high magnetic field gradients it cannot be used for light storage. One way around this is to use a conservative trap, such as a red-detuned dipole trap introduced in Sec. 9.2.3.

This trapping field must be π -polarised and detuned on the order of Δ_{fs} or higher so as not to affect the different m_F splitting, and therefore the gradient, along or across the ensemble. Achieving π polarisation requires the addition of a constant dc magnetic field to create a quantisation axis. Such a beam, tightly focused and collimated, should provide sufficient trapping in the radial direction but much weaker trapping along the ensemble, as $F_d = -\nabla(U_d) = -\bar{U}_d \nabla(I(r, z))$ and $\partial_r I \gg \partial_z I$. The radius of the ensemble will then become approximately equal to the waist of the trapping laser w_d , which can be reduced down to the order of 10 μm . This is desirable for our system as a reduction in area for the acS-gradient-creating laser ($\mathcal{A}_{acS} \geq 2Lw_d$) will lead to an increase in the intensity, and therefore bandwidth, for a given power (see Sec. 10.2.3). Realistic experimental parameters give L on the order of one centimetre [397, 398]. To obtain the longest storage times possible the dipole laser should be detuned far from resonance. Assuming an easily obtainable experimental wavelength of $\lambda_d = 1064$ nm and a waist of $w_d = 10$ μm , a trap of depth $U_d = 1$ mK would be achievable with a power of 1.5 W, giving a maximum scattering rate of $\Gamma_d = 4$ s⁻¹ ¹.

As discussed in the previous chapter, to avoid longitudinal diffusion and its associated

¹Calculations included the counter-rotating terms ($\omega_o + \omega_l$) as the rotating-wave approximation is not valid for large detunings where $\omega_o - \omega_l \ll \omega_o + \omega_l$ no longer holds.




Figure 10.4: Scattering Rate for a One Megahertz ac Stark Gradient. Scattering rate shown for (i)→(v) $m_F = -2 \rightarrow m_F = 2$ states of the $F = 2$ level, as a function of detuning $\Delta_{1/2,2}/2\pi$ (solid lines, left-hand scale). Also shown is the intensity of the acS field required to produce the 1 MHz splitting (dashed line, right-hand scale). These were calculated using Eq. 10.1 and 9.12 without the rotating-wave approximation.

frequency blurring, it is important to be able to also trap atoms in the z -direction. One solution would be to split the trapping field, using a 50:50 non-polarising beam-splitter, and send it into the ensemble from both directions to create a standing-wave trap. Trapping maxima would occur at $\lambda_d/2$, with the difference in detuning between two sites being

$$\Delta\delta_t = \frac{2\pi\mathcal{B}_s\lambda_d}{2L}, \quad (10.8)$$

assuming a linear intensity profile. As discussed in the previous chapter, storing light in an optical lattice such as this could provide storage times of hundreds of milliseconds [389] to seconds [119]. The probe and coupling fields could easily be combined with the trapping field using dichroic mirrors with a cut-off wavelength of approximately 1000 nm.

One issue that could arise with light being sent into a periodic medium, such as atoms in a standing-wave dipole trap, is the possibility of the light being reflected or dispersed by the medium. However, work by J. Nunn et al. showed that, as long as the periodicity of the lattice is not close to half the probe wavelength λ_p , there would be no adverse effects [406]. As here we would have $\lambda_d/\lambda_p = 1.34$, we expect no issues with using a standing-wave dipole trap for our memory.

10.2.3 Gradient Creation

Wavelength Selection

A critical parameter that must be determined is the wavelength of the acS laser to be used, as this will set a limit on the maximum frequency splitting possible for a given laser power and intensity distribution, as well as the scattering rate of the system. To optimise the wavelength we must balance the desired behaviour (i.e., frequency splitting) with the undesired effect of light scattering by the atoms determined by $\bar{\Gamma}_{F,m_F}$. It will be

assumed that the acS field will be circularly polarised to maximise $\bar{\delta}_t$ for large detunings (see Sec. 10.2.1).

From Eq. 10.4 we can see that the splitting, and therefore bandwidth, depends not only on $1/\Delta$ but also on the ratio of $\Delta_{1/2,F}/\Delta_{3/2,F}$. As the detuning becomes large compared to Δ_{fs} , $\Delta_{1/2,F}/\Delta_{3/2,F} \rightarrow 1$ as $\Delta_{3/2,F} = \Delta_{1/2,F} + \Delta_{fs}$, and therefore the bracketed term in Eq. 10.4 heads to zero. This means that $\bar{\delta}_F$ will head towards zero faster than $1/\Delta$ for large detunings. This can be seen from the difference in Fig. 10.3(b) between $\bar{\delta}_t$ and the dashed line, which shows the path $\bar{\delta}_t$ would take if it had only a $1/\Delta$ dependence. This deviation occurs on the order of terahertz detuning and is important as it leads to Fig. 10.4. This shows the scattering rates of different m_F states of the $F = 2$ level for an acS bandwidth of 1 MHz, assuming the same splitting arrangement used for Fig. 10.3(b) and $q = 1$, i.e.,

$$\Gamma_{2,m_F}(\Delta, 1) = \frac{10^6 \cdot \bar{\Gamma}_{2,m_F}(\Delta, 1)}{\bar{\delta}_t(\Delta, 1)}, \quad (10.9)$$

with the corresponding intensity profile also shown².

From this figure it can be seen that the behaviour for the $m_F = 0, -1, -2$ levels is similar to that described for laser trapping in Sec. 10.2.1, i.e., decreasing scattering rate with increasing detuning for a set trap depth. However, for $m_F = 1$ a minimum for the scattering rate is present at approximately 5 THz, with the scattering rate flattening out until 20 THz when it starts to decrease again. The $m_F = 2$ state shows even more peculiar behaviour, with a maximum appearing at approximately 20 THz. A positive slope indicates that $\bar{\delta}_t$ is decreasing faster than $\bar{\Gamma}$. This is to do, in part, with the dependence of $\bar{\delta}_t$ on $\Delta_{1/2,F}/\Delta_{3/2,F}$ discussed above, as well as the relative strengths of $A_{3/2,g_i}$ and $A_{1/2,g_i}$ and the levels which are allowed to contribute to $\bar{\Gamma}$. All scattering rates converge at detunings much larger than Δ_{fs} . The scattering rate for the $F = 1$ states are approximately the same in magnitude but for the opposite m_F state, with the $m_F = -1$ ground state also containing this minima. In both cases the unusual behaviour occurs for the states which are raised in energy with respect to $m_F = 0$.

Due to the small probe approximation (see Ch. 3), most of the population will remain in the $F = 1$ state during the storage and retrieval processes. It would therefore be advantageous to make use of the minima and the $F = 1, m_F = -1$ state as, though lower scattering rates can be achieved at much larger detunings, the laser intensities required become impractical. For instance, to achieve the same, or smaller, scattering rate than at the minima (approximately 11 s^{-1}) requires a detuning greater than 40 THz. At this detuning it requires nearly ten times the laser intensity to achieve the same bandwidth as at the minima. For the minima detuning $\bar{\delta}_1 \simeq \bar{\delta}_2 \simeq 50 \text{ Hz W}^{-1} \text{ cm}^2$ and $\bar{\delta}_{1,2} \ll \bar{\delta}_F$.

Beam Shaping

The optimal intensity distribution - that which will give the maximum bandwidth for a given power P_o - depends firstly on the orientation of the ensemble. As discussed in Sec. 10.2.2, we envisage a cylindrical ensemble with $L = 1 \text{ cm}$ and $R = w_d = 10 \text{ }\mu\text{m}$, as shown in Fig. 10.1(b).

One of the simplest intensity profiles to produce is a focused Gaussian, where

$$I_G(r', y) = I_o \text{Exp} \left[-2r'^2/w_o^2 \right], \quad (10.10)$$

²Previous footnote on rotating-wave approximation also applies here.

with w_o the waist of the beam occurring at position $y = 0$ (the centre of the ensemble, assuming the acS field is propagating along the y -axis), I_o is the maximum intensity and r' is the radial component in the x, z -plane. The only constraint on the acS beam shape is that it is monotonic along the ensemble to avoid reabsorption issues (see Sec. 3.4.2). In the case of a Gaussian beam, this means that over half the power will be lost, limiting the total splitting possible for a given laser power, with $I_o = 2P_o/(\pi w_o^2)$.

To use the largest fraction of the beam possible, we can set $w_o = 2L/3$ so that if the centre of the beam is at $z = 0$, approximately 99% of the remaining intensity will fall on the ensemble. Figure 10.1(c)(i) shows this intensity profile along and across the ensemble for unit power. As the gradient η is not linear, different frequency components will be stored with different efficiencies, as can be seen from Eq. 3.54. In the case described above, where $R \ll L$, the frequency change along the x -direction for a set z -location will be negligible.

If one of the experimentally simplest intensity profiles for the acS beam is the one described above, then one of the most efficient will be one that covers only the ensemble, decreasing linearly from maximum intensity I_o at one end to zero at the other and with no change in intensity in the x -direction. Such an intensity profile will be of the form

$$\begin{aligned} I_L(r', y) &= I_o (1 - z/L), |x| \leq R \\ &= 0, |x| > R \end{aligned} \quad (10.11)$$

over the length of the ensemble $0 \leq z \leq L$. In this case $I_o = P_o/(LR)$ and therefore the maximum intensity, and bandwidth, achievable will be over 500 times larger than a Gaussian beam with the same power due to the smaller area the beam occupies. This can be seen from Fig. 10.1(c)(ii), which shows the intensity profile I_L for unit power along and across the ensemble.

To change an initially Gaussian beam to one with an intensity profile like that from Eq. 10.11 requires a beam shaper. These devices (for instance deformable mirrors, phase plates or liquid crystal spatial-light modulators - SLMs) can be highly efficient ($\varepsilon > 0.9$) and can be used to create nearly any desired beam shape with resolution on the order of 1000×1000 pixels for SLMs. This not only provides us with a method for optimising the acS laser intensity profile, but would also allow for spectral manipulation of the pulse to be carried out with the ability to produce complex gradients and switching arrangements, potentially with better resolution than those demonstrated in Ch. 7.

10.2.4 Switching Protocols

There are two components that make up the switching protocol: which m_F states within the two F levels are to be used; and the method for switching the acS field. These will be discussed in turn.

Selecting m_F Levels

Selection rules determine which m_F states of the two different levels can be used, depending on the polarisations of the probe and coupling fields (q_p and q_c respectively), with the total change in angular momentum being given by

$$\Delta m_F = m_2 - m_1 = q_p - q_c. \quad (10.12)$$

As $m_1 = -1$ has already been decided upon as state $|1\rangle$ (see Sec. 10.2.3), the above equation can be rearranged to give

$$m_2 = -1 + q_p - q_c. \quad (10.13)$$

We can determine the probe and coupling polarisations that will produce this maximum splitting by substituting the above relationship into Eq. 10.6 to give

$$\begin{aligned} \bar{\delta}_t &= \bar{\delta}_{1,2} + q(\bar{\delta}_1 - (-1 + q_p - q_c)\bar{\delta}_2) \\ &\simeq q\bar{\delta}_F(2 + q_c - q_p), \end{aligned} \quad (10.14)$$

with the approximations $\bar{\delta}_1 \simeq \bar{\delta}_2 \gg \bar{\delta}_{1,2}$ and $|q| \neq 0$ discussed previously. This reveals that $\bar{\delta}_t$ can range from 0 to $3\bar{\delta}_F$ (remembering $|m_2| \leq 2$) with the latter being possible for either $q_c = 1, q_p = 0$ or $q_c = 0, q_p = -1$.

Using the previous value of $\bar{\delta}_F$ determined in Sec. 10.2.3, this would give a system bandwidth of $\mathcal{B}_s = 150 \text{ Hz W}^{-1} \text{ cm}^2$. If we combine this value with the maximum intensity achievable with the intensity profile I_L we find $\mathcal{B}_s = 150 \text{ kHz W}^{-1}$ and therefore system bandwidths on the order of 1 MHz would be obtainable with less than 10 W of acS laser power using the optimal detuning. Also, using one of the above optimal level schemes with $\delta_t = 3\bar{\delta}_F$, the scattering rate per megahertz splitting given in Eq. 10.9, which was calculated using $\delta_t = 2\bar{\delta}_F$, would be reduced as the intensity required to reach the same bandwidth would be reduced by a factor of 2/3.

The polarisations for the probe and coupling fields mean that they could not be combined or separated using polarisation optics. To detect only the probe at the end of the memory a frequency selective measurement, such as heterodyne detection, or spatial filtering with a small angle, could be used.

Field Switching Method

To cause rephasing of the atomic dipoles we must be able to invert the detunings of the atoms. If we make the same assumption as in previous sections, i.e., $|q| = 1$, two switching methods become readily apparent.

The first method involves reversing the intensity profile along the ensemble $I(z) \rightarrow I(L - z)$. This is equivalent to the magnetic field switching method used in previous chapters in that, by reversing the intensity profile about the centre of the trap ($z = L/2$), the detunings are also reversed about this point, i.e., $\delta_t(z) \rightarrow \delta_t(L - z)$. The intensity profiles for this method of switching are shown in Fig. 10.5(a)(i), with the corresponding frequency gradients shown in Fig. 10.5(a)(ii). This process involves no change in frequency of the stored pulse with respect to the input pulse.

The second method involves switching the polarisation of the field $q \rightarrow -q$ while keeping the same intensity gradient, as shown in Fig. 10.5(b)(i). This is a slightly more complicated process as the detunings are no longer reversed around the centre of the ensemble, with $\delta_t(z) \rightarrow -\delta_t(z)$. This method still results in an echo being produced. However, the stored pulse will now be frequency-shifted with respect to the input pulse, as can be seen from the corresponding frequency gradients shown in Fig. 10.5(b)(ii). This is because a switch from $\delta_t(z) \rightarrow -\delta_t(z)$ is equivalent to a switch about the centre and an offset of $\delta_{t_{max}}$ being added. In a three-level system this frequency shift can be overcome by altering the coupling field frequency in such a way as to cancel the initial shift. It

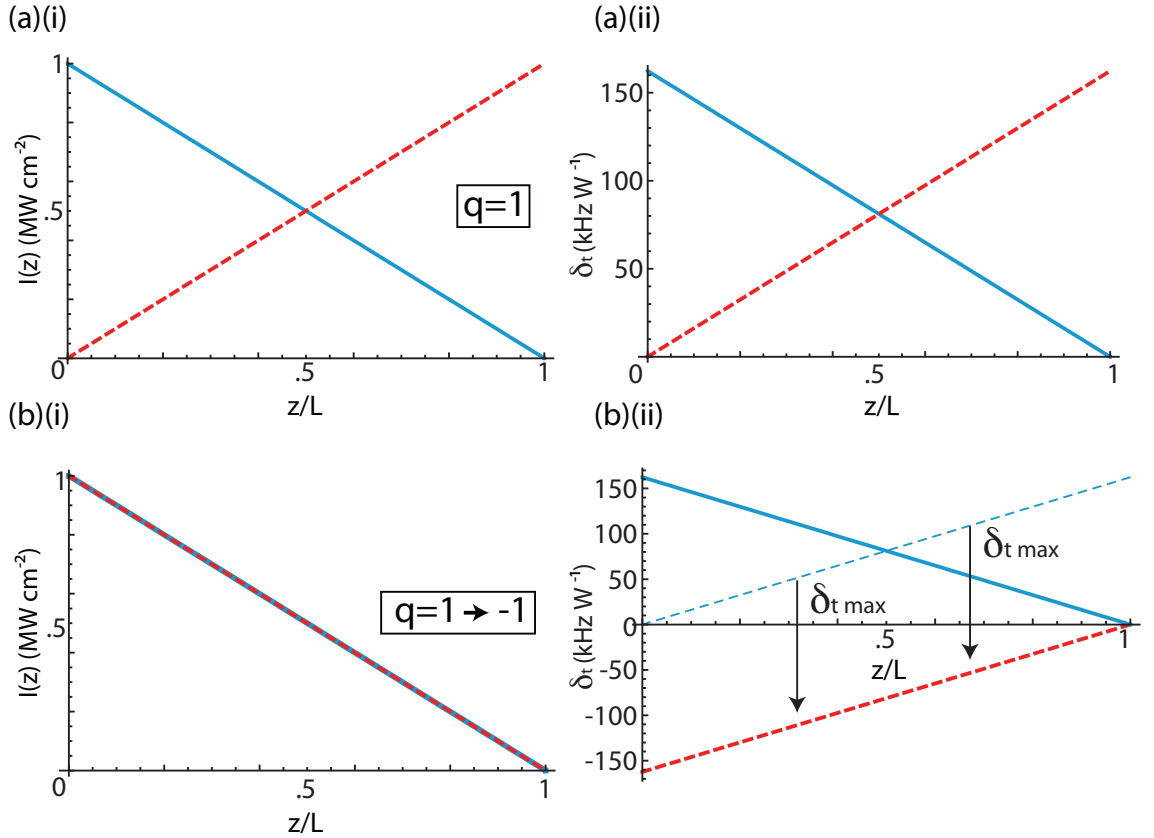


Figure 10.5: ac Stark Switching Methods. (a)(i) Intensity profile per unit power and (a)(ii) corresponding frequency splitting along the ensemble for Switching Method 1, with initial (solid line) and final (dashed line) gradients shown. Here the polarisation of the acS field $q = 1$ for both initial and final gradients. (b)(i) Intensity profile per unit power and (b)(ii) corresponding frequency splitting for Switching Method 2, with initial (solid line) and final (dashed line) gradients shown. Here the initial polarisation of the acS field is $q = 1$, while the final polarisation is $q = -1$. This causes a frequency offset to the gradient of $-\delta_{t,max}$. Both (a) and (b) use the linear intensity profile I_L (see Eq. 10.11) and the maximum $\bar{\delta}_t \simeq 3\bar{\delta}_F$.

should also be noted that $\bar{\delta}_{1,2}$ will not contribute to $\bar{\delta}_t$ for this switching method as it only depends on the detuning and intensity of the acS laser, both of which will be constant.

This second switching method would suggest itself as the easiest to implement as all that is required to switch η would be a Pockels cell, which have switching times down to nanoseconds. It would not, however, allow for different frequency gradients and therefore filtering or manipulation of the pulse. Apart from this, if the second switching method were to be used, then the minima for the scattering rate found in Sec. 10.2.3 will only apply for either the read or write stages of the memory. This is because, if the polarisation is reversed, the scattering behaviour of the levels will be reversed (i.e., $\bar{\Gamma}_{2,1}(q = -1) \rightarrow \bar{\Gamma}_{2,-1}(q = 1)$).

The first switching method would allow for different gradients but involves much longer switching times (on the order of milliseconds for LCSLMs). The combination of beam shapers (BSh) and Pockels cells (PC) shown in Fig. 10.1(a) would allow for flexibility in beam shaping and fast switching times. If no spatial filtering is desired then




Figure 10.6: ac Stark Cold Atom GEM Efficiency. (a) Total efficiency of the Λ -GEM memory as a function of pulse length t_p for storage times of $t_s = t_p$ and different ratios of $|\Omega_c/\Delta_{1p}| =$ (i) 0.01 (blue), (ii) 0.003 (red), and (iii) 0.001 (green). (b) Total efficiency for storage of pulses of length $t_p = 20 \mu\text{s}$ as a function of normalised storage time t_s/t_p for (i) single pulse (blue); and (ii) multiple pulse storage (red), with $|\Omega_c/\Delta_{1p}| = 0.02$. For all traces $\Delta_{1p} = -2\pi \times 2 \text{ GHz}$, $\Delta_{ac} = -2\pi \times 5 \text{ THz}$, $\mathcal{B}_s = \mathcal{B}_G = 9\sqrt{2}/(\pi t_p)$, $q_c = 1$, $q_p = 0$, $g = 2\pi \times 1.5 \text{ MHz}$, $N = 2.5 \times 10^6$ atoms, $\epsilon_L = 0.4$, $L = 1 \text{ cm}$ and $R = 10 \mu\text{m}$. NCR denotes the non-classical regime for the memory for large photon numbers (shaded).

only PC2 and BSh1 would be needed, with the beam shaper determining the shape of the gradient and the polarisation switch causing the rephasing of the atoms. To allow different gradients to be used, an extra Pockels cell PC1 and beam shaper BSh2 could be used. In this case, the second gradient could be prepared in advance and PC1 used to select which beam shaper to use. The acousto-optic modulator (AOM) would be used to switch the acS beam on or off to decrease the scattering rate due to this field (see following section). Alternatively another PC and polarising beam-splitter could be used.

10.3 Limiting Factors

10.3.1 Time Scales

As discussed in the previous chapter, one main advantage of moving from warm to cold atoms is the extended storage times that can be achieved. There are two timescale of importance to our memory: (i) the trap lifetime τ_{trap} ; and (ii) the coherence time τ_{coh} .

The trap lifetime depends on both the scattering rate of the trapping laser, as well as the rate of interatomic collisions. These will also affect the coherence time. Using the trap parameters from Sec. 10.2.2 gives a coherence time of approximately $1/\Gamma_d = 250 \text{ ms}$. Coupling the trap scattering rate with the recoil energy per emission $E_{rec} = (\hbar k)^2/2m$ [378], where m is the mass of rubidium and k is the wave-vector of the transition (here taken to be for the D2 transition), give the trap lifetime for a given depth U_d to be

$$\tau_{trap} = \frac{mU_d}{\hbar^2 k^2 \Gamma_d}. \quad (10.15)$$

Using this equation and U_d and Γ_d determined in Sec. 10.2.2, we find that the trap lifetime

will be on the order of ten seconds and should therefore not affect the coherence time of the system (see below).

Another effect that must be considered is the inelastic collision rate between the atoms in the trap, as this will also cause a loss of coherence as well as trap population if the atoms are not all situated in the $F = 1$ level. The collisional loss rate is described in Eq. 9.9. Here we will ignore the α term, as it will be much smaller than β , and take $\beta = 5 \times 10^{-11} \text{ cm}^3 \text{ s}^{-1}$. For the densities we are expecting ($n \approx 10^{11} \text{ atoms cm}^{-3}$), and assuming a constant density within the trapping volume for simplicity, Eq. 9.9 will give an initial, and maximum, collision rate of $\beta n \simeq 30 \text{ s}^{-1}$.

Apart from the two sources of decoherence mentioned above, which will also affect the trap lifetime, there are two others that must be considered, namely the scattering rate due to the acS field as well as the coupling field. The scattering rate for the acS field can be determined as a function of system bandwidth \mathcal{B}_s , similarly to Eq. 10.9, to be

$$\Gamma_{ac}(\Delta) = \frac{\bar{\Gamma}_{1,-1}(\Delta, 1)}{\bar{\delta}_t(\Delta)} \mathcal{B}_s. \quad (10.16)$$

For the optimal detuning of $\Delta_{1/2,2}/2\pi = 5 \text{ THz}$ discovered in Sec. 10.2.3 for the $F = 1, m_F = -1$ state with $q = 1$, this simplifies to $\Gamma_{ac} \simeq 7 \cdot 10^{-6} \mathcal{B}_s$ for $\bar{\delta}_t \simeq 3\bar{\delta}_F$.

The optical depth of the system depends on the Rabi frequency of the coupling field Ω_c and the one-photon detuning Δ_{1p} , as was shown in Eq. 3.54. It therefore makes sense to express the effect of the coupling field on the ground state in terms of these two parameters by using the relation between coupling field intensity I_c and Rabi frequency $I_c = 2\hbar^2 \epsilon_0 c \Omega_c^2 / d_{2,2'}^2$ [256] to give

$$\Gamma_c(\Delta_{1p}, \Omega_c) = \frac{2\hbar^2 \epsilon_0 c}{d_{2,2'}^2} \bar{\Gamma}_{1,-1}(\Delta_{1p} - \Delta_{hfs}, q_c) \Omega_c^2, \quad (10.17)$$

$d_{2,2'}$ being the D1 dipole transition strength between the $F = 2$ and $F' = 2$ levels. For pulses much longer than $1 \mu\text{s}$ and $|\Omega_c/\Delta_{1p}| \geq 0.001$, Γ_c will be much greater than Γ_{ac} .

The total scattering rate for the system is simply the sum of the individual rates. The acS and coupling fields are only needed during the reading and writing phases of the memory process, which each last for a minimum period determined by the pulse length (i.e., $t_{r/w} \geq t_p$ for single pulses, but becoming longer for multiple-pulse storage). If we define a background scattering rate Γ_{bg} to include all the decoherence effects that are constantly present, i.e., scattering from the trapping laser, collisions and loss from the trap, and a read/write scattering rate of $\Gamma_{rw} = \Gamma_{ac} + \Gamma_c$, then we can determine the total storage efficiency to be

$$\epsilon_s(t_p, \tau) \leq \text{Exp}[-2t_p \Gamma_{rw}] \text{Exp}[-(2t_p + t_s) \Gamma_{bg}], \quad (10.18)$$

where t_s is the time the pulse is stored in the memory.

10.3.2 Efficiency

There are three main experimental factors that will affect the storage efficiency of this system: the number of atoms N initially present in the MOT; the loading efficiency from the MOT to the dipole trap ϵ_L ; and the size of the gradient which is applied $|\eta(z)| =$

$2\pi\mathcal{B}_s/L$. Substituting these values into Eq.s 3.42 and 3.54 gives the read/write efficiency of the system

$$\varepsilon_{rw} = \left(1 - \text{Exp} \left[-\frac{g^2 \epsilon_L N L}{c \mathcal{B}_s} \left(\frac{\Omega_c}{\Delta_{pc}} \right)^2 \right] \right)^2. \quad (10.19)$$

To investigate the maximum efficiency as a function of storage time we assumed, to maximise optical depth and minimise scattering, that the bandwidth of the system was equal to the bandwidth of the pulse, i.e., $\mathcal{B}_s = \mathcal{B}_p$. For Gaussian pulses we can define the pulse bandwidth to be $\mathcal{B}_G = 9\sqrt{2}/(\pi t_p)$, assuming 99% of the intensity field is stored in the memory. Figure 10.6(a) shows the total efficiency $\varepsilon_t = \varepsilon_{rw} \cdot \varepsilon_s$ for different values of Ω_c/Δ_{1p} , with storage time $t_s = t_p$, and using $\mathcal{B}_s = \mathcal{B}_G$.

As can be seen from the figure, for each ratio of Ω_c/Δ_{1p} there was a maximum efficiency. This was a combination of two effects: firstly the increase in optical depth as the pulse length increases due to the smaller bandwidths required, and extra decoherence that will occur as the storage time increases. The decoherence is greater for larger values of Ω_c/Δ_{1p} due to the increase in Γ_{rw} this entails, and this can also be seen from Fig. 10.6(a). In this regime, however, the main source of decoherence was the atomic collisions, with an approximate coherence time of 30 ms.

For short pulses ($t_p \ll 1$ ms) and large values of $|\Omega_c/\Delta_{1p}| \simeq 0.02$ efficiencies approached unity as optical depths were high and the background decoherence effects would be negligible. This was therefore a good regime to investigate the delay-bandwidth product (DBP) of the memory. Here we define $\text{DBP} \equiv t_s/t_p$ as this will give an indication of the number of pulses (or bits) that can be stored in the memory at one time (see Sec.s 1.4.5 and 3.2.2).

Figure 10.6(b) shows ε_t for (i) single- and (ii) multiple-pulse storage with $t_p = 20$ μ s. The difference between the two is that, for multiple-pulse storage, the coupling and acS fields are left on at all times, while this is not the case for single-pulse storage. As can be seen, single-pulse storage with efficiency greater than 90% (dashed line) occurs up to approximately $t_s = 130t_p$. For multiple-pulse storage, this drops to 50 pulses due to the extra coupling field scattering, while 350 can be stored above the classical efficiency limit of $\varepsilon_t = 0.5$. To date, a maximum of twenty pulses have been simultaneously stored using Λ -GEM with $\varepsilon_t < 0.2$ [4], and the largest DBP with total efficiency above 50% is 74 using cold atom electromagnetically-induced transparency³ [13]. To store 20- μ s-long pulses requires a system bandwidth of approximately 200 kHz, and therefore an acS laser power of less than 2 W in the optimised regime.

10.4 Experimental ac Stark Gradient Generation

Moving from a theoretical proposal to an experimental demonstration can hold many traps for the unwary. We therefore decided to test out the ac Stark gradient creation experimentally. The critical component of this experiment was the choice of laser. The criteria we had to satisfy were: wavelength as close to the optimal value of 806 nm (i.e., 5 THz detuned); power as high as possible; and a source that could be focused down to as small a spot as possible. In the end we settled on a Coherent fibre array package (FAP) system, with a wavelength of 808 nm, power output of 30 W and a minimum spot size of $d = 450$ μ m.

Without ready access to specialised beam shaping equipment such as a deformable mirror or SLM, we attempted to create a linear gradient using a 25.4 mm² top-hat profile

³Note: in Ref. [13] they define pulse width as the full-width-half-maximum.

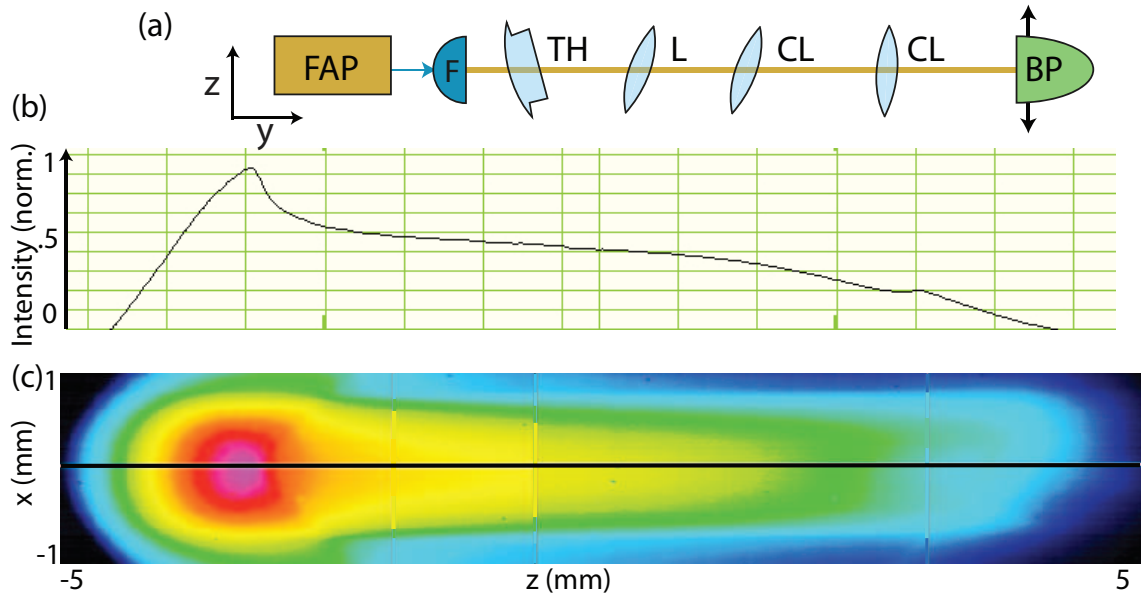


Figure 10.7: ac Stark Gradient Generation. (a) Schematic diagram of set-up used to create acS gradient, with: FAP - fibre array package laser; F - fibre out-coupler; TH - top hat profile; L - circular lens; CL - cylindrical lens; BP - translating beam profiler. (b) 1D beam-profiler-generated plot of intensity. (c) 2D beam-profiler-generated plot of intensity. These tests were conducted at low power ($\ll 1$ W).

from Holo/or. Simulations by the company indicated that a uniform intensity gradient could be created for the FAP by misplacing it (i.e., a slight y -offset in Fig. 10.7(a)), with a minimum size in the x -direction of $200\ \mu\text{m}$ for an optimised top-hat profile with a working distance of 35 cm, at over 95% efficiency. However, with the profile we had, and the fact that the FAP was a highly divergent light source, we were not able to achieve this. To improve the linearity and compactness of our ac Stark gradient we used a variety of lenses (both circular and cylindrical) at various angles, along with the top-hat profile. The power out of the FAP was set very low ($\ll 1$ W) so that we could use a beam profiler to directly detect the output intensity profile of the laser. This set-up is shown in Fig. 10.7(a).

Figures 10.7(b) and (c) show the results of this test. Over about 60% of the approximately 1 cm length the gradient was at least monotonic if not entirely linear. While the width of the gradient in the y -direction was approximately 2 mm. Taking a $6 \times 2\ \text{mm}^2$ area would give us a bandwidth of only 37.5 kHz, assuming no loss and that 80% of the laser power lies in the middle 6 mm. The main difference from this low figure and the $150\ \text{kHz W}^{-1}$ determined in Sec. 10.2.4 is the assumed width of the beam (2 mm versus $10\ \mu\text{m}$). Even with the smallest possible beam width determined for the Holo/or profile of $200\ \mu\text{m}$, the bandwidth would only be 214 kHz for a 1-cm-long memory. It would therefore be worth investigating the use of SLMs and deformable mirrors to see if they could improve both the linearity and the width of the gradient further.

10.5 Conclusions

In this chapter we have proposed a new method of gradient creation using the ac Stark effect, to go along with our move to a cold atom memory. We investigated the experimental viability of this proposal. Firstly, the trapping mechanism was investigated and it was determined that a far-detuned standing-wave dipole trap would not interfere with the acS gradient and have a lifetime on the order of seconds, while providing a small area for the acS laser. It was found that there is an optimal detuning for the acS field of approximately 5 THz to minimise scattering and maximise bandwidth. With this detuning, and an optimised beam shape and level scheme, a bandwidth of 150 kHz W^{-1} could theoretically be created. Different methods for gradient switching were investigated and a scheme devised that would allow a switch between any two arbitrary gradients on the order of nanoseconds using Pockels cells and spatial light modulators.

Factors that would limit the system such as scattering due to the trapping, coupling and acS lasers, as well as collisions between atoms were examined. These were combined with the read/write efficiency to model the total system efficiency as a function of storage time and it was seen that for long pulses ($t_p \gg 1 \text{ ms}$), the coherence time was limited by the atomic collisions to be on the order of tens of milliseconds. It was also found that for short pulses ($t_p \ll 1 \text{ ms}$), efficiencies approach unity and 50 pulses could be stored in the memory at one time with efficiency greater than 90%. The acS laser power required to produce the necessary memory bandwidth to store these pulses would be less than 2 W in the optimised regime.

Finally, we very briefly attempted to experimentally create an ac Stark gradient using a high power fibre laser system. Without the aid of specialised beam shaping equipment we were only able to create a monotonic gradient of approximately 6 mm in length, with a total bandwidth of less than 37.5 kHz. This could be improved by at least an order of magnitude by reducing the acS beam's vertical width with more specialised equipment.

We therefore conclude that using an ac Stark gradient with cold atoms is an experimentally viable option for improving the Λ -GEM system in terms of gradient creation, switching and manipulation, as well as storage times and time-bandwidth products achievable using cold atoms.

Experiments Towards Cold Atom GEM

Turk: “Come on, man, learn by doing! Learn by doing.”
 - “My First Day” *Scrubs*, NBC, 2 October 2001, TV Series.

11.1 Introduction

The move from warm atom storage towards cold atom storage brings with it many layers of experimental complexity. In this chapter we present two experiments exploring this cold new world, with the aim being to identify issues that need addressing, as well as to devise the experimental set-up and develop the techniques required to produce the best possible results for Λ -GEM using cold atoms. We use EIT as a diagnostic tool for Λ -GEM as they are similar in a couple of key areas: namely they both use three-levels in a Λ configuration; and they both use a strong coupling field, while EIT does not require magnetic field switching and can therefore be implemented with less equipment.

The remainder of the chapter is divided into two main sections: Sec. 11.2 presents work carried out in a three-tier ^{87}Rb magneto-optical trap (MOT) and Sec. 11.3 presents work carried out on a Cs MOT.

Part of the work presented in Sec. 11.3 has been published in the journal article:

- L. Giner, L. Veissier, B. M. Sparkes, A. S. Sheremet, A. Nicolas, O. S. Mishina, M. Scherman, S. Burks, I. Shomroni, D. V. Kupriyanov, P. K. Lam, E. Giacobino, and J. Laurat, *Experimental Investigation of the Transition between Autler-Townes Splitting and Electromagnetically-Induced Transparency*, Physical Review A **87**, 013823 (2013).

11.2 Experiment 1 - Three-Tier Rubidium MOT and Dipole Trap

The first experimental set-up that was used to investigate cold atom light storage was a three-tier ^{87}Rb MOT, designed as a gravimeter [407]. The set-up was divided into three different sections: the *memory set-up* used to create the EIT fields required (Sec. 11.2.1); the *atom trapping set-up* that produced the cold atomic ensemble used to store the EIT probe (Sec. 11.2.2); and the *interface and detection set-up* where the two set-ups were combined and data recorded. The same storage and atom trapping set-ups were used

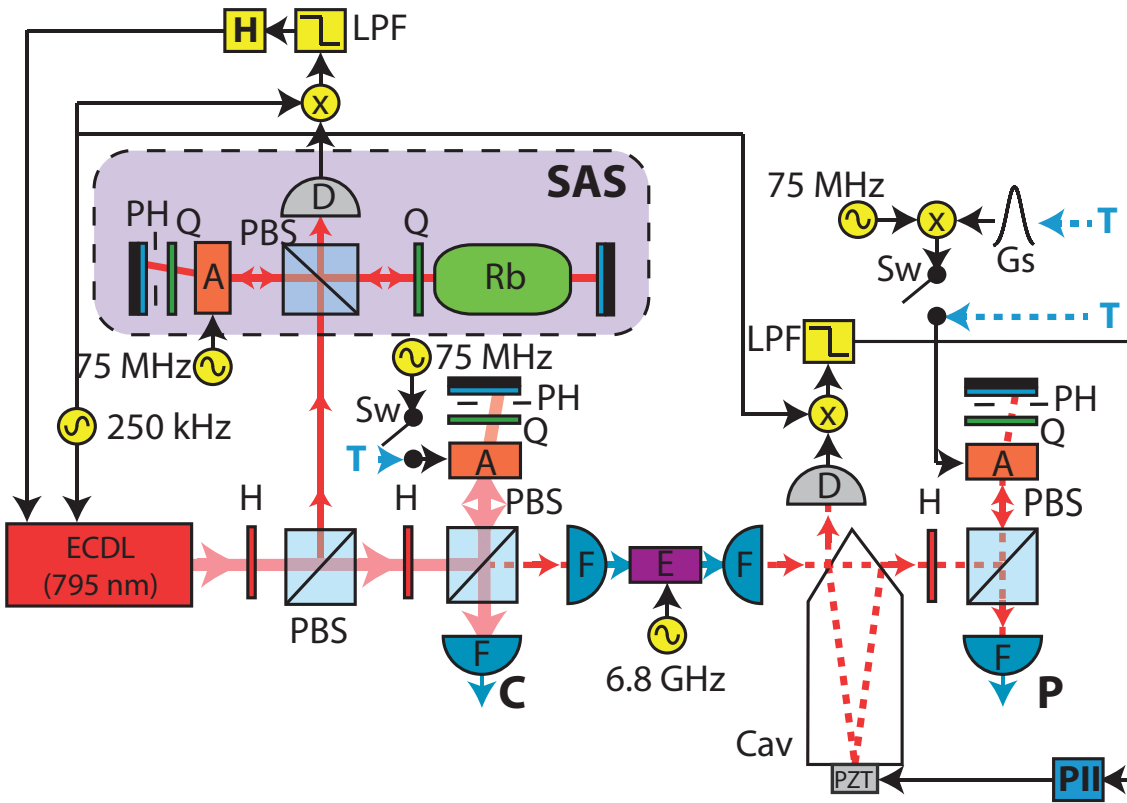


Figure 11.1: Cold Atom Memory Set-Up, Expt. 1, with: ECDL - 795 nm external cavity diode laser; C - coupling field (thick red line); P - probe field (dashed red line); SAS - saturated absorption spectroscopy set-up (shaded region); H - half wave-plate; PBS - polarising beam-splitter; Q - quarter wave-plate; A - acousto-optic modulator; E - electro-optic modulator; Cav - ring cavity; PZT - piezo-electric transducer; PH - pin-hole; F - fibre coupler; D - detector; Rb - natural mixture Rb vapour cell; Sw - switch; LPF - low-pass filter; H - MOGLabs controller; PII - LabVIEW-based controller (see Ch. 4); Gs - Gaussian envelope function; T - digital trigger signal.)

for all sub-experiments included in this section. The interface and detection set-ups were altered to try to optimise the storage and these will be presented in Sec.s 11.2.4-11.2.7, along with the various results obtained.

11.2.1 Memory Set-Up

The memory fields set-up for Expt. 1 was similar to that used in Ch.s 6-8, and is shown in Fig. 11.1. Unlike previous chapters, however, the laser used was not a titanium:sapphire but a MOGLabs external cavity diode laser (ECDL) with a centre wavelength of approximately 795 nm, with 100 mW power output. Also unlike previous chapters, the low temperatures of the cold atoms meant that there was minimal Doppler broadening of the Rb absorption line and, therefore, it was important that both the probe and coupling fields were as close to the actual Rb resonances as possible for EIT experiments. To overcome this issue, the ECDL frequency was locked to an atomic absorption line using a saturated absorption spectroscopy (SAS) set-up (shaded region in Fig. 11.1). A modulation of 250 kHz was placed on the laser current and a small portion ($< 1\%$) of the laser power was double-passed through a natural mixture Rb vapour cell. By passing light

through the gas cell from both directions, sub-Doppler linewidth transition features could be resolved. Some SAS spectra are shown in App. A, for more information on SAS see Ref. [204]. Error signals for the atomic transition peaks were obtained by mixing down the detected current-modulated signal, similar to the methods used in App. C, inside the Moglabs controller. These error signals were used to lock the laser to the desired atomic transition.

Approximately 8 mW of the 100 mW was sent to the coupling field set-up, and the rest was sent to a fibre-coupled electro-optic modulator (EOM). This was used in combination with a ring cavity¹, as was the case in the warm atom Λ -GEM chapters, to allow the flexibility to have the probe address the $F = 1 \rightarrow F'$ transition and the coupling field address the $F = 2 \rightarrow F'$ transition, or vica verse, by driving the EOM at 6.8 GHz. The ring cavity was used to select the probe sideband to transmit. It should be noted that, due to the non-Gaussian profile of the beam leaving the ECDL, the coupling efficiency to the fibre was very low ($< 20\%$) and only 3 mW was obtained at the output of the cavity.

Acousto-optic modulators (AOMs) were again used to allow for fine frequency control, as well as temporal control and pulsing. However, unlike previous experiments, the probe and coupling fields were double-passed through the AOMs. This was to try to minimise the change in coupling efficiency of the light into the fibres that would be experienced if the position of the laser (dependent on the AOM modulation frequency) were to be altered [211]. The AOMs used had an optimum efficiency at a driving frequency of approximately 80 MHz, and therefore both the probe and coupling fields were shifted by nearly 160 MHz with respect to the initial laser by the time they entered the fibre. As mentioned above, it was important that the probe and coupling fields be able to be locked to an atomic transition. To counter this issue, a third AOM was included in the SAS arm of the experiment. All AOMs shifted light in the same direction while being driven at the same frequency and, therefore, when the SAS set-up was locked, all fields were on resonance with the desired atomic transition.

The probe and coupling fields were then fibre-coupled to a second table. The powers of the probe and coupling fields out of their fibres were 400 μ W and 2 mW respectively.

11.2.2 Atom Trapping Set-Up

As mentioned above, the atom trapping set-up used for Expt. 1 was originally designed as an experiment to measure gravity, with an experimental sequence aimed at transferring atoms to the top MOT, which was then dropped vertically for approximately 30 cm. It was therefore overly complicated for our needs. However, one should never pass up an opportunity, graciously provided by Dr. Nick Robins and the ANU Quantum Sensors and Atom Laser Group, and the system still had all the necessary components for a test-bed quantum memory experiment. Here, only the details of the set-up relevant for our atom trapping and light storage applications will be discussed. For further detail on this set-up see Ref. [407]. Fig. 12.14 also shows a detailed version of a cooling and trapping set-up very similar to this.

A schematic of the overall MOT set-up is shown in Fig. 11.2(a). Three MOGLabs ECDLs were used to derive all light fields shown in the figure. One ECDL, locked to the ^{87}Rb D2 $F = 2 \rightarrow 1, 3$ crossover using a SAS set-up like the one shown in Fig. 11.1 (again see figures in App A). One difference between this SAS lock and the one presented in the previous section was that here the locking modulation was not provided by the

¹This cavity was locked using the digital locking code described in Ch. 4.

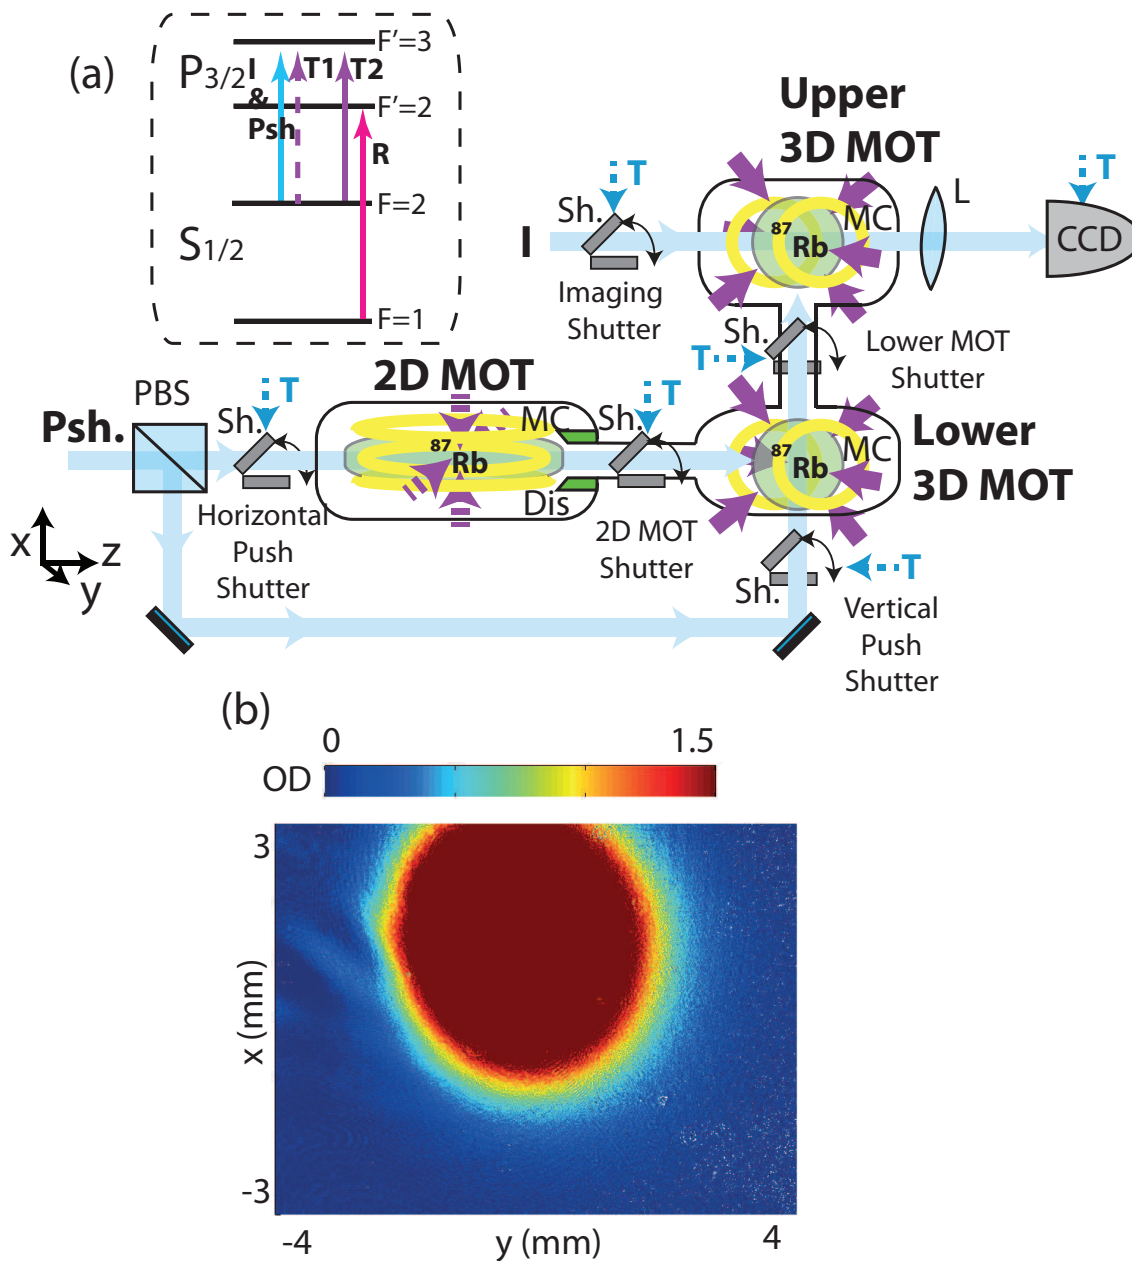


Figure 11.2: Cold Atom Trapping Set-Up, Expt. 1. (a) A schematic diagram of the relevant components of the atom trapping set-up used for Experiment 1, with: Psh. - push field (blue); I - imaging field (also blue); T1/2 - trapping fields 1 (2D MOT, dashed purple) and 2 (3D MOTs, solid purple); R - repump (pink); PBS - polarising beam-splitter; Sh. - shutter; Dis - Rb dispenser; ^{87}Rb - trapped ^{87}Rb atoms; L - lens; CCD - charged-coupled device camera; T - digital triggers; MC - MOT coils (electronics behind switching on/off not shown). Inset shows energy level diagram, the repump is not shown in the main figure but was co-propagating with the trapping field. (b) Absorption image of the upper 3D MOT used for the memory experiments, colour bar shows OD scale.

laser current, but by a set of coils placed around the reference Rb cell which modulated the internal atomic states via the Zeeman shift (see Sec. A). This is known as Zeeman modulation [408], and is more desirable than current modulation locking for certain applications as there is no modulation present on the laser. A double-pass AOM driven at approximately 106 MHz was then used to up-shift the frequency of the laser so that it was just below resonance with the D2 $F = 2 \rightarrow F' = 3$ closed transition. This field was used for the 3D trapping (with approximately 30 mW of power for each 3D MOTs), imaging and push beams. A second ECDL was similarly locked to the D2 $F = 2 \rightarrow 2, 3$ crossover and another double-pass AOM used to lock it to the same transition as the 3D MOT laser. However, this AOM was placed in the locking loop to allow for maximum power, with approximately 100 mW reaching the atoms. This means that the AOM could not be used to gate this 2D MOT trapping laser, only the shutters. Finally, a third ECDL was locked to the D2 $F = 1 \rightarrow 2, 3$ crossover, and then a single-pass 80 MHz AOM used to shift the laser to be resonant with the D2 $F = 1 \rightarrow F' = 2$ transition. This laser was used as the repump field (see Sec. 9.2.1). These fields were then fibre-coupled, and linked to a second table. The final frequencies of the MOT fields are shown in Fig. 11.2(a).

The second table contained the three atom traps, shown in Fig. 11.2(a). Firstly, there was a 2D MOT, which in our experiment was always kept on. As discussed in Sec. 9.4.2, this used only four trapping beams and a quadrupole magnetic field to constantly trap atoms, released from an Alvatec natural mixture Rb dispenser, in the y - and x -directions. Atoms were transferred from the 2D MOT to the lower 3D MOT (i.e. a “normal” MOT created using six trapping beams and a pair of anti-Helmholtz coils) via a horizontal push beam that applied radiation pressure to the atoms in the z -direction. Atoms were moved from the lower MOT to the upper 3D MOT via a vertical push beam. The repump field was combined with the trapping fields before the MOTs (not shown in figure) to keep the atomic population in the $F = 2$ level. The imaging beam was used to probe the upper MOT to determine its shape and how many atoms were present in the $F = 2$ level (and therefore the optical depth - OD). The second table was surrounded by six rectangular compensation coils, used to cancel the dc magnetic field at the upper 3D MOT. All shutters, switches and AOMs (both amplitude and frequency) were controlled via a LabVIEW code purpose-built for cold atom experiments using National Instruments (NI) PXI-6733 (analog) and PXI-6534 (digital) cards. The code minimised memory storage on the analog and digital cards by using a timing array to notify the cards when the next change of values would occur, allowing for minute-long runs with timing precision down to $2 \mu\text{s}$.

The actions required as part of a normal experimental run are shown in Tab. 11.1. For the first 35 s of the run the horizontal push beam was turned on to load the lower MOT from the 2D MOT. For the next 7 s the vertical push beam shutter was turned on to start loading the upper MOT from the lower MOT, which also continued to fill. The vertical push beam was then turned off and for the next 500 ms the current in the lower MOT coils was slowly ramped down to zero (i.e., no magnetic field) to avoid the Eddy currents a fast switch off could produce. At this time ($t = 42.5$ s) all loading was stopped by turning off the horizontal push beam. Also starting at 42.5 s, the frequency of the upper MOT trapping fields and the current of the upper MOT coils were ramped down to zero during 50 ms. This was to allow for polarisation-gradient cooling [357] of the atomic cloud (see Sec. 9.2.2). From here, the repump and trapping fields were turned off/on at various times, the memory fields were triggered (by sending a trigger signal from the digital NI PXI-6534 cards controlling the triggering to the PXI chassis controlling the memory set-up via field

Time (s)	Action
Initially	All coils on All shutters closed All beam amplitudes (except probe and control) on
0.00	Horizontal push shutter opened 2D MOT shutter opened lower MOT shutter opened control and probe off
35.00	Vertical push shutter opened
42.00	Vertical push shutter closed Lower MOT coil current ramped down to 0 A (500 ms)
42.50	Horizontal push shutter closed Lower MOT shutter closed Push/Imaging beam amplitude off Upper MOT trapping beam frequency ramped down to 0 (50 ms) Upper MOT coil current ramped down to 0 (over 50 ms)

Table 11.1: Main Run Times for Expt. 1. This is the main part of the runs used in Expt. 1, showing the time and what action was taken at that time

programmable gate array - FPGA - cards). These variations to the cold atom run are discussed in the following sections, after a brief detour via absorption imaging.

11.2.3 Absorption Imaging

To be able to optimise something, one invariably has to be able to quantify certain aspects of it. In this case, we would like to know how many atoms are present, and the density of the ensemble (i.e., the OD). Fluorescence imaging, where the scattered photons from the MOT are detected, can prove useful for measuring MOT loading rates but is fraught with danger for the uninitiated when it comes to optimising total atom number, and gives no information on density.

One way of characterising a MOT is through absorption imaging. Absorption imaging allows us to measure the OD of a MOT using Beer's law

$$\text{OD} \simeq \ln[I_t/I_o], \quad (11.1)$$

where I_t is the transmitted imaging beam intensity after passing through the MOT, and I_o is the intensity measured without the MOT present. If we use a charged-coupled device (CCD) camera then we can calculate the OD from Eq. 11.1 for each pixel using a program, such as Matlab, to produce an absorption image of the MOT, giving us the required information. If, however, the atomic sample is very dense, then no light will get through the MOT on resonance and therefore not much information can be provided. One way to counter this is to move the imaging field away from resonance. As the absorption of light by atoms away from resonance will follow a Lorentzian decay, we can image the MOT far from resonance to provide all required detail, and then scale the measured OD

Time (s)	Action
Initially	Probe/coupling field on (50 nW)
42.5519	Upper MOT trapping beam amplitude off Repump beam off
42.553	Imaging shutter open
42.5531	Push/Imaging beam amplitude on CCD Camera trigger
42.5532	Push/Imaging beam amplitude off
42.56	Push/Imaging shutter open
42.5601	Push/Imaging beam amplitude on CCD Camera trigger
42.5602	Push/Imaging shutter closed

Table 11.2: Beam Alignment Run for Expt. 1. The remainder of the run (following from Tab. 11.1) for aligning probe and control fields to the MOT.

to on-resonance OD using [409]

$$\text{OD}_{res} = \frac{\Delta^2 + \gamma^2/4}{\gamma^2/4} \text{OD}_{\Delta}, \quad (11.2)$$

where Δ is the detuning of the imaging beam and γ is the excited-state decay rate of atoms in question. This equation assumes that the imaging field intensity is very small, otherwise the corrected linewidth $\gamma' = \gamma\sqrt{1 + I/I_{sat}}$, I_{sat} being the saturation intensity.

In practise, to image the entire MOT the imaging beam must be made larger than the atomic cloud. Also, after the MOT, an imaging system must be put in place to image a picture of the MOT onto the CCD camera. This is achieved by placing a lens between the MOT camera and focusing using the simple lens equation

$$\frac{1}{f} = \frac{1}{l_o} + \frac{1}{l_i}, \quad (11.3)$$

where f is the focal length of the lens, l_o is the distance from the MOT to the lens, and l_i is the distance from the lens to the camera. The magnification of the object is given by the ratio of l_i/l_o .

Figure 11.2(b) shows an image of a MOT produced using the above method and extra steps shown in Tab. 11.2, with the probe and coupling fields blocked. The image that produced I_t was taken 1.2 ms after the MOT fields are turned off. The reference image that produced I_o was taken 7 ms later, after all atoms had left the MOT chamber. The image beam was on resonance with the cloud (i.e., the D2 $F = 2 \rightarrow F' = 3$ transition) and the magnification was 1:1. As can be seen, the average atomic cloud produced had a radius of approximately 2-3 mm and a peak optical depth, determined by fitting a 2D Gaussian to the data, of between 3.5-5. The number of atoms present varied between $1\text{-}2 \cdot 10^8$.

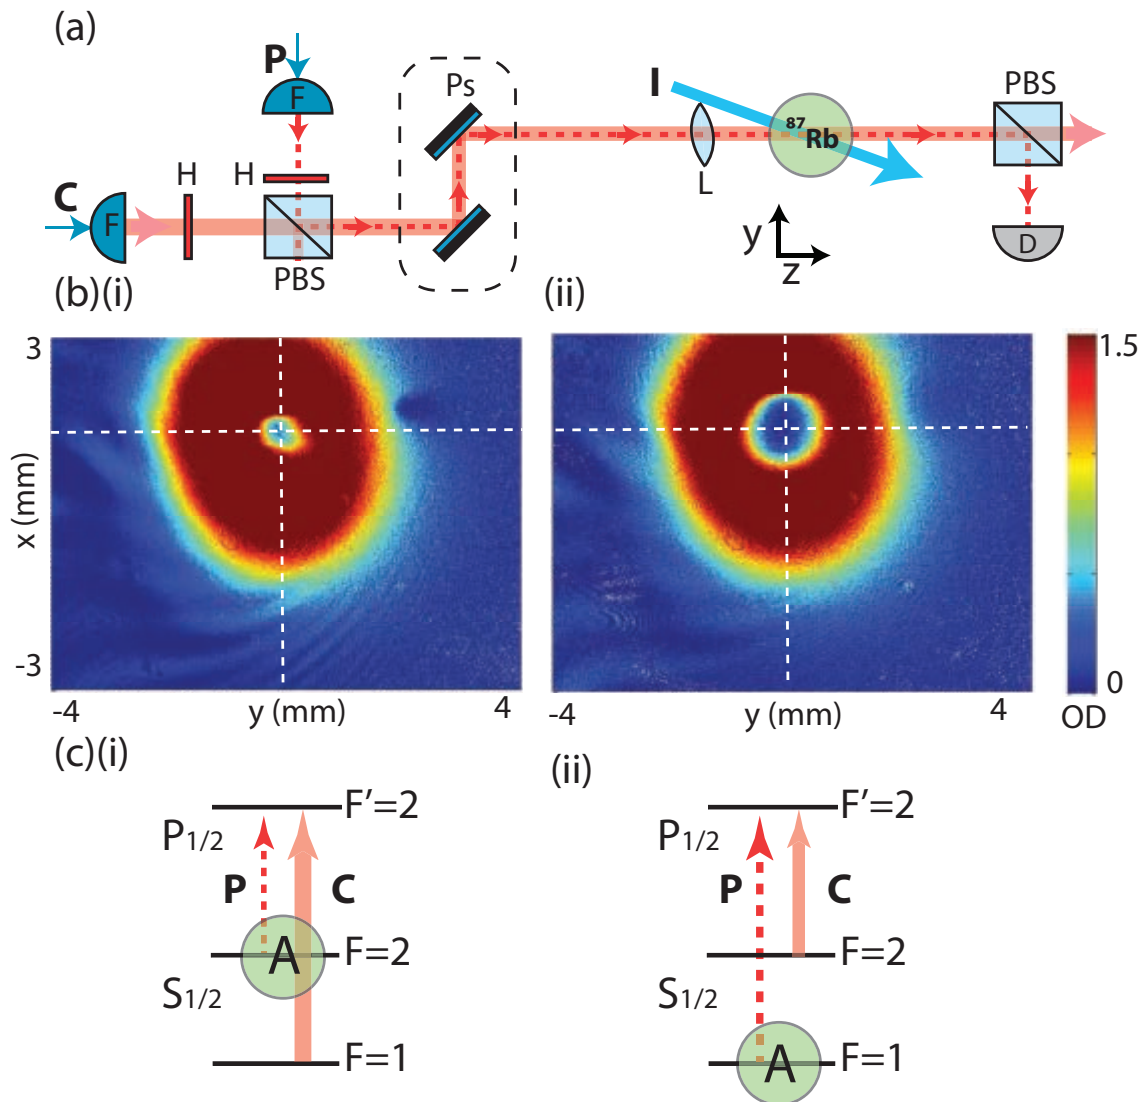


Figure 11.3: Interface and Detection Set-up 1a,b (Opposite Polarizations). (a) Schematic diagram of set-up, with: C - coupling field (thick red line); P - probe field (dashed red line); I - imaging field (blue line); F - fibre out-coupler; H - half wave-plate; PBS - polarising beam-splitter; Ps - periscope; L - lens D - detector; ^{87}Rb - trapped ^{87}Rb atoms (MOT coils and trapping fields not shown). (b) Absorption images used for alignment of (i) the probe and (ii) the coupling fields through the MOT. (c) The atomic population (A), energy levels, and probe and coupling transitions used for interface and detection set-up for (i) 1a and (ii) 1b.

11.2.4 Interface and Detection Set-up 1a - Opposite Polarisation, $F = 2$

Befitting the first attempt at combining the light storage and atom trapping experiments, the first interface set-up was the easiest to implement. As shown in Fig. 11.3(a), the probe and coupling fields had opposite linear polarisations exiting the fibres. The two fields were combined on a polarising beam-splitter (PBS) and sent, by way of a periscope, through the upper 3D MOT. A 150 mm focal length lens was used to focus the probe and coupling fields into the MOT, ensuring that they interacted only with the central atoms. After they passed through the MOT, the two fields were separated on another PBS. The probe field intensity was then measured using the same type of photo-detector used for the warm atom measurements. The output from the photo-detector was sent through a Stanford Instruments pre-amplifier before being viewed on an oscilloscope.

The imaging set-up was used to initially align the probe and coupling fields through the MOT, following the beam alignment run detailed in Tab. 11.2. In this run the probe (coupling) field was left on for the entire run with minimal power (approximately 50 nW) while the coupling (probe) field was blocked. Both fields were on resonance with the D1 $F = 2 \rightarrow F' = 2$ transition. This was achieved by locking the SAS set-up to the aforementioned transition, while not driving the fibre-coupled EOM so that all the probe power was left in the carrier. For this run the upper MOT trapping and repump field amplitudes were turned off at the same time so that atoms remained in the $F = 2$ level. Too much power in the probe or coupling fields would lead to all atoms being ejected from the MOT. To achieve the low powers (tens of nanowatts) required to only eject those atoms in the direct path of the beams, and therefore aid with the alignment of the beams, a neutral-density (ND) filter was placed in front of the beams leaving the fibres. As can be seen from Fig. 11.3(b), the probe and coupling beams were aligned so as to pass through the centre of the MOT and the coupling beam was made larger than the probe.

It was initially decided to place the probe field on the D1 $F = 2 \rightarrow F' = 2$ transition (and therefore the coupling field on the D1 $F = 1 \rightarrow F' = 2$ transition), as the population of atoms inside the MOT was initially in the $F = 2$ level. This arrangement of population, and probe and coupling field transitions is illustrated in Fig. 11.3(c)(i), with the run being described in Tab. 11.3(a)(i). The probe and coupling fields were again gated via electronic switches added before their respective AOMs, and triggered by transistor-to-transistor logic (TTL) signals from the same LabVIEW code that ran the atom trapping experiment. An additional TTL signal, again from this code, was used to trigger a 1 ms frequency scan of either the probe or coupling field's frequency through an Agilent waveform generator. Using this method the EIT transmission feature was investigated and the results shown in Fig. 11.4.

Firstly from this figure it can be seen that there was nearly 100% absorption of the probe field. However, with the coupling field on, and with relatively high powers for this situation, the amount of enhanced transmission was minimal: only between 10-20%, with a full-width-half-maximum (FWHM) of approximately 0.5-1.0 MHz.

11.2.5 Interface and Detection Set-up 1b - Opposite Polarisation, $F = 1$

As the enhanced transmission results from Expt. 1a were approximately an order of magnitude below the near-unity transmission expected from a cold atomic ensemble, it was decided to work in a similar regime to many other groups that have investigated EIT using cold atoms (see Sec. 9.4.1) and move the atomic population to the $F = 1$ level and place the probe on the D1 $F = 1 \rightarrow F' = 2$ transition. This arrangement of population, probe

Time (s)	Action
42.55	(a) Repump amplitude left on (b) Repump amplitude off
42.5519	Upper MOT trapping beam amplitude off (a) Repump amplitude off Coupling field on
42.553	(i) Probe field on (i) Trigger 1 - for scanning probe/control frequency (1 ms) (ii) Trigger 2 - for probe pulse
42.554	Coupling field off (b) Repump amplitude on
42.5555	(b) Repump amplitude off
42.56	Imaging shutter open
42.561	Push/Imaging beam amplitude on CCD Camera trigger
42.5611	Push/Imaging beam amplitude off
42.56	Push/Imaging shutter open
42.5601	Push/Imaging beam amplitude on CCD Camera trigger
42.5602	Push/Imaging shutter closed

Table 11.3: Light Storage Runs for Expt.1a-d. The remainder of the run (following from Tab. 11.1) for the EIT light storage protocols attempted for: (a) Expt. 1a (probe on the D1 $F = 2 \rightarrow F' = 2$ transition); (b) Expt. 1b,1c (probe on the D1 $F = 1 \rightarrow F' = 2$ transition); (i) to investigate EIT transmission feature; (ii) to investigate pulse absorption/delay.




Figure 11.4: Scanning EIT in MOT on $F = 2$, Expt. 1a (Opposite Polarizations). EIT transmission feature obtained with the probe field on the D1 $F = 2 \rightarrow F' = 2$ transition with a probe power of 250 nW and coupling field powers of (i) 250 μ W, (ii) 400 μ W and (iii) zero. Two averages per trace.

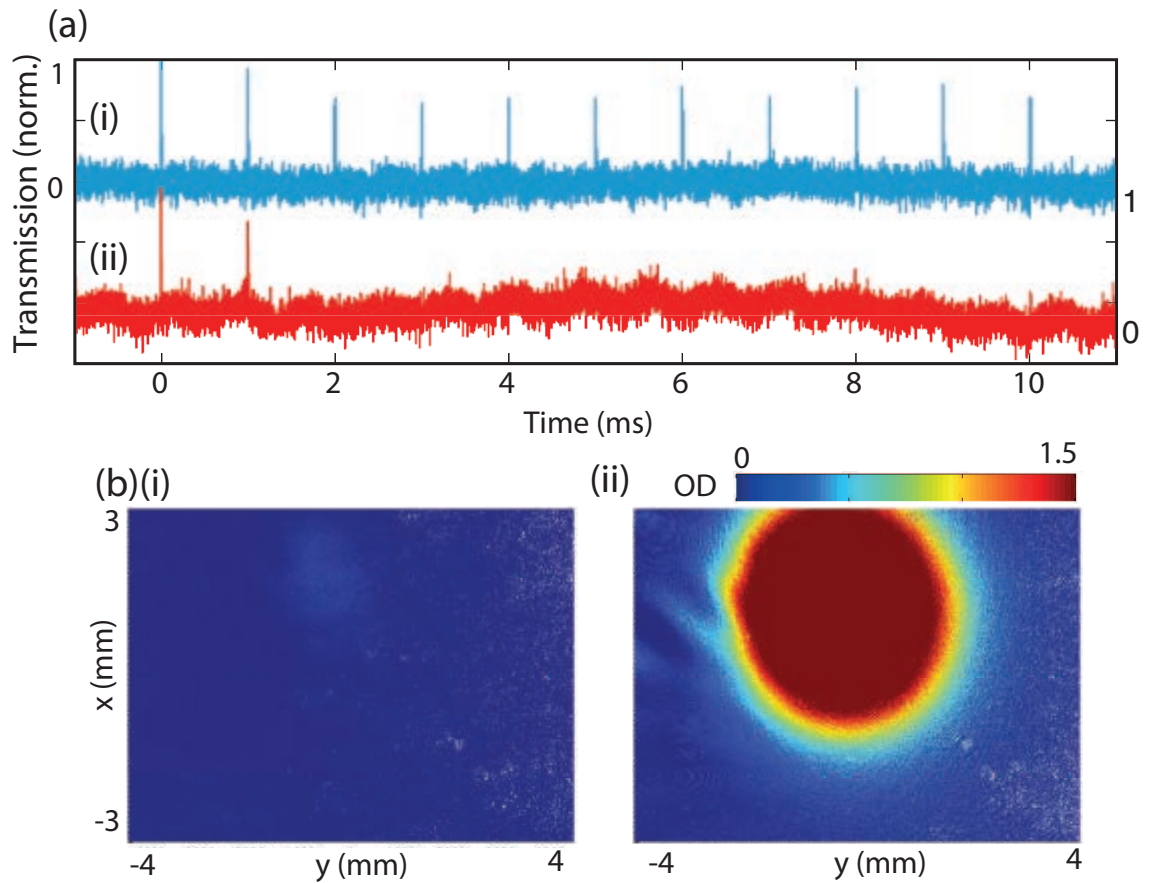


Figure 11.5: Population Pumping from $F = 2 \rightarrow F = 1$ Level, Expt. 1b (Opposite Polarizations). (a) Probing of the atomic population in the $F = 1$ ground state with probe pulses with $\approx 4 \mu\text{W}$ on the D1 $F = 1 \rightarrow F' = 2$ transition for (i) no MOT (blue, left-hand scale) and (ii) MOT present (red, right-hand scale). Eight averages per trace. (b) Absorption images for (i) run described in Tab. 11.3(b) and (ii) after an additional 1.5 ms repump pulse. Colour bar shows OD scale.




Figure 11.6: EIT in MOT on $F = 1$, Expt. 1b (Opposite Polarizations). (a) The EIT feature obtained with the probe field on the D1 $F = 1 \rightarrow F' = 2$ transition with approximately $1 \mu\text{W}$ probe power. Coupling field power shown on figure. Four averages per trace. (b) Absorption and delay of probe pulses with a maximum probe peak height of approximately $1 \mu\text{W}$. Coupling field powers shown on key. Four averages per trace.

and corresponding coupling field transitions is illustrated in Fig. 11.3(c)(ii). To achieve this reshuffling of population from the $F = 2$ level, the run for Expt. 1b differed from the previous run, with the repump field being turned off approximately 2 ms before the MOT trapping field (see Tab. 11.3(b)).

Two methods were used to investigate this population reshuffle. Firstly, a series of probe pulses on the D1 $F = 1 \rightarrow F = 2$ transition were sent into the cold atomic ensemble in 1 ms increments around this population pumping event (i.e., from 42.5500 s to 42.5605 s in the run), with the repump being turned off 1 ms later than indicated in Tab. 11.3) at 42.551 s. The absorption of these pulses was then investigated. This data is shown in Fig. 11.5(a). As can be seen, approximately 3 ms after the first pulse there was enough population in the $F = 1$ level to completely absorb the probe pulses. To investigate this further the MOT was imaged at this time, shown in Fig. 11.5(b)(i). As can be seen, no atoms were detected in the $F = 2$ level². To ensure that the atoms had indeed been pumped to the $F = 1$ level, the repump field was turned back on for 1.5 ms after the light storage experiment was complete, but before the CCD camera was triggered, to pump the population back from the $F = 1 \rightarrow F = 2$ level. One such image is shown in Fig. 11.5(b)(ii), illustrating that the atomic population was not being lost, but simply pumped to the other level, as desired.

With the probe on the D1 $F = 1 \rightarrow F' = 2$ transition (and the coupling field on the D1 $F = 2 \rightarrow F' = 2$ transition) the EIT feature was again investigated for different coupling field powers using the run described above. The results from this investigation are shown in Fig. 11.6(a). As can be seen, the enhanced transmission feature reached a maximum of approximately 60% for a coupling field power of $C_p = 450 \mu\text{W}$ with a FWHM of approximately 2 MHz. This seems quite wide when compared to the work of M. Arikawa et al. who, using a similar system, produced an EIT feature with a FWHM of approximately 600 kHz [131]. It is also interesting that as the power decreased the top of the transmission feature flattened off. This was most pronounced for the lowest power

²Remembering that the imaging beam was on resonance with the D2 $F = 2 \rightarrow F' = 3$ transition.

traces (coupling power of 5-10 μW) which gave the impression that the EIT feature may have been being broadened. One final point to note from this figure is that, as the coupling power was decreased, the absorption of the probe decreased. This was because the atoms absorbed the probe light and moved to the $F = 2$ level and, without the coupling field to pump them back to $F = 1$, they remained there. This was an illustration of the sensitivity of the cold atomic ensemble to even very small continuous probe powers, as opposed to the relative insensitivity of warm vapour cells.

Having investigated the EIT feature, the next step was to investigate the enhanced transmission and delay of pulses sent through the MOT. This was achieved using the run described in Tab. 11.3(b)(ii), which involved pulsing the probe using an external Picotest arbitrary waveform generator (with a FWHM of approximately 1 μs), which was again triggered by a TTL signal from the LabVIEW atom trapping code, and mixed with the signal from the waveform generator with the probe frequency signal (approximately 75 MHz) before an electronic switch (see Fig. 11.1). The results of this are shown in Fig. 11.6(b). As can be seen, without the coupling field nearly all the input pulse is absorbed and, with the coupling field on, a maximum transmission of approximately 55% for a coupling power of greater than 150 μW was obtained. This agreed with the maximum pulse transmission data shown in Fig. 11.6(a), where a maximum transmission of approximately 60% was seen. The other important parameter that could be gleaned from this trace was the delay of the pulse through the atomic ensemble compared to the pulse when there were no atoms. The delay is related to the inverse of the EIT feature width seen in Fig. 11.6(a) and therefore will decrease as the coupling field power increases and feature broadens (see Sec. 3.3.1). This can be seen to be the case, however the maximum delay was only approximately 150 ns for a coupling power of 5-10 μW . This again agrees with the EIT feature as the delay $\approx L/v_p = \sqrt{\text{OD}}/(2\pi \cdot \text{FWHM})$, where L is the length of the memory (see Eq.s 3.20 and 3.22). The FWHM of the EIT feature for small coupling powers was still approximately 1 MHz due to unusual broadened shape of the feature discussed above.

11.2.6 Interface and Detection Set-up 1c - Same Polarisation, $F = 1$

As the 150 ns delay obtained above was too short to be able to store a pulse efficiently on the timescales and pulse widths we were working with, the next step was to try the same experiment but with the same polarisations for the probe and coupling fields. This meant that a PBS could no longer be used to combine and separate the two fields. Instead, similar to many other papers dealing with EIT in cold atomic ensembles, the probe and coupling fields were set-up to travel along different paths but overlap inside the ensemble. This set-up is shown in Fig. 11.7(a). Again the probe and coupling fields were aligned through the MOT using the imaging system (see Sec. 11.2.4). A quarter wave-plate was added to the set-up to allow the polarisations to be altered to right or left circular.

The EIT feature was again investigated using the run described in Tab. 11.3(b)(i) for both left and right circular polarisations, as well as linear. The results from this investigation are shown in Fig. 11.7(b). These traces were taken with relatively low coupling field power (approximately 5-10 μW) to remove any power broadening effects (see Sec. 5.3.4). It is interesting to see that with one circular polarisation the transmission feature appeared on one side of the two-photon resonance (i.e., $\delta = 0$) while for the other orthogonal circular polarisation it appeared on the other side. For the same linear polarisation, as was the case for orthogonal linear polarisations for low coupling powers, the feature appeared broadened between these two. This was a clear indication that there was a non-zero mag-

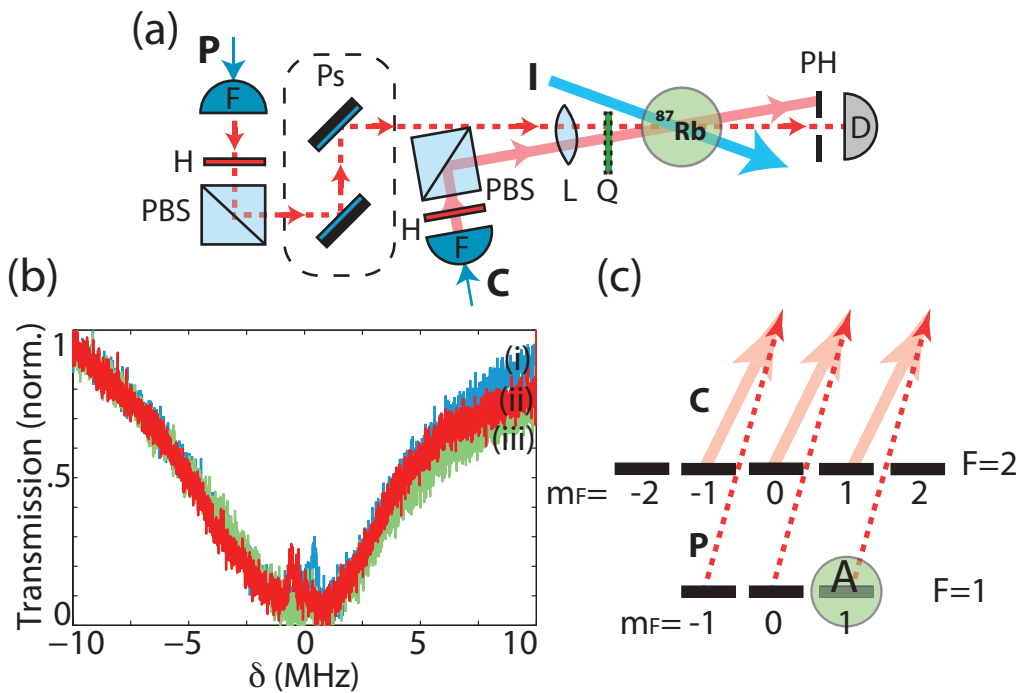


Figure 11.7: Interface and Detection Set-up 1c (Same Polarisations), and Optical Pumping Effects. (a) Schematic diagram of set-up, with: C - coupling field (thick red line); P - probe field (dashed red line); I - imaging field (blue line); F - fibre out-coupler; H - half wave-plate; PBS - polarising beam-splitter; Q - quarter wave-plate (removable); Ps - periscope; L - lens; PH - pin-hole D - detector; ^{87}Rb - trapped ^{87}Rb atoms (MOT coils and trapping fields not shown). (b) EIT transmission window obtained with the probe field on the D1 $F = 1 \rightarrow F' = 2$ with approximately $1 \mu\text{W}$ of power and $5\text{-}10 \mu\text{W}$ coupling power for: (i) both same-circularly polarised (+); (ii) both same-circularly polarised (-); and (c) both same-linearly polarised. Four average per trace. (c) An illustration of optical pumping within the two ground states of rubidium with probe and coupling fields with the same circular polarisation, A indicates the position of the atomic population.

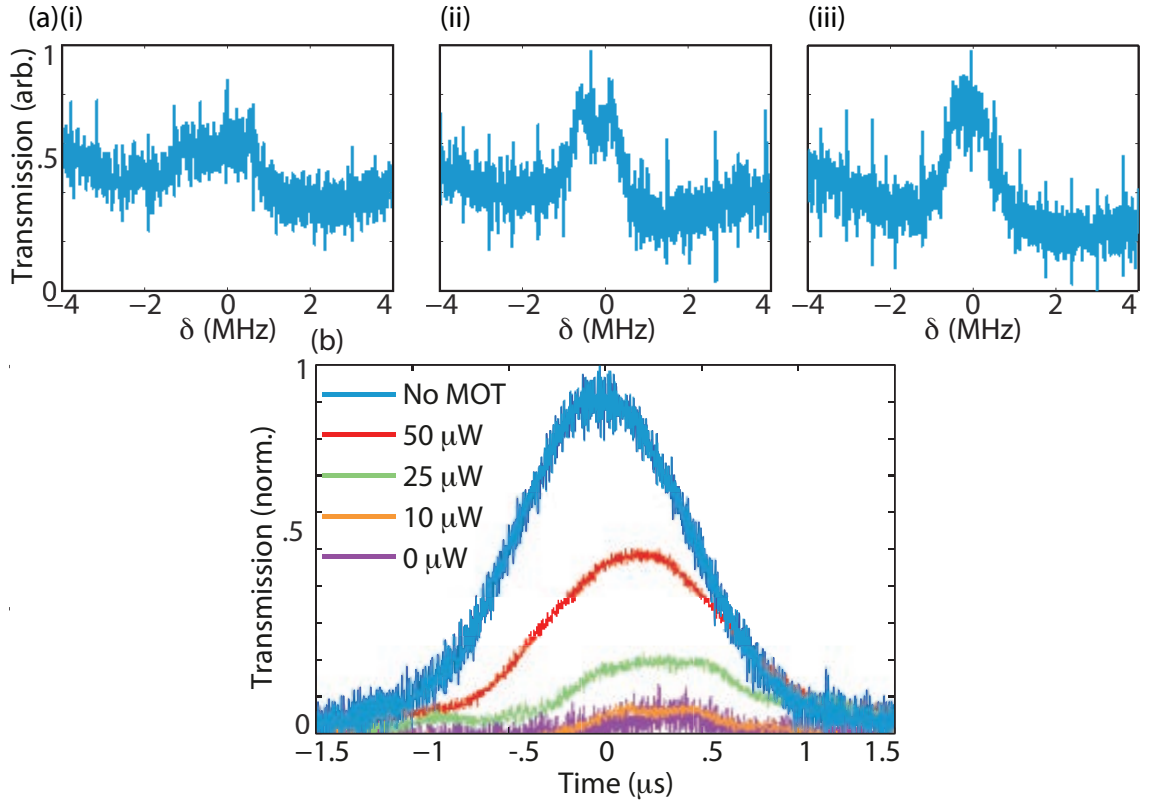


Figure 11.8: Optimising Magnetic Field Compensation, Expt. 1c (Same Polarisations). (a) EIT features with the probe on the D1 $F = 1 \rightarrow F' = 2$ transition. (i)→(iii) denote the optimisation of the currents in the compensation coils. Probe power was approximately $1 \mu\text{W}$, coupling power was $40 \mu\text{W}$. Two averages per trace. (b) Absorption and delay of probe pulses with a maximum probe peak height of approximately $1 \mu\text{W}$. Coupling field power shown in key. Four averages per trace.

netic field across the ensemble as, in the presence of such a field, the different m_F states will be at different energies. The effect of having the same circular polarisation was then to pump atoms into either the $\pm 1 m_F$ state, depending on the direction - right or left - of the circular polarisation. This is illustrated in Fig. 11.7(c). If all m_F states had the same energy, which would be the case if $|B| = 0$, then the EIT feature would appear in the same location for both right and left circular polarisations. The broadening of the feature for linear polarisations could also be explained this way, as linear polarisation is a mixture of right and left circular polarisation and so the atoms spread out among the three different m_F states at different energies, broadening the line.

We attempted to reduce the dc magnetic field to zero across the ensemble by altering the currents in the three sets of compensation coils. The results of this, using linear polarisations for the probe and coupling fields, are shown in Fig. 11.8(a) from (i) the initial compensation coil currents through to (iii) the optimal currents, where there was no longer any difference between circular and linear polarisations. However, the FWHM of the EIT feature was still on the order of 1 MHz for a coupling field power of $40 \mu\text{W}$.

With this optimisation complete, more pulses were sent in to the cold atomic ensemble to again test the delay and transmission of the system. The results from this investigation




Figure 11.9: Raman Absorption, Expt. 1c (Same Polarizations). (a) Raman feature for different one-photon detunings (see key) and constant coupling power ($500 \mu\text{W}$). (b) Raman feature for set one-photon detuning (-15 MHz) and varying coupling power (see key). Probe power was approximately $4.5 \mu\text{W}$ on the D1 $F = 1 \rightarrow F' = 2$ transition, four averages per trace.

are shown in Fig. 11.8(b). As can be seen, the maximum delay was increased from 160 ns to approximately 250 ns. Unfortunately this was still not long enough to see storage of light in the system using EIT. The small delays achieved reflect the low optical depth: compared to the work described in Ref. [13] with an OD of 160, an optical depth of 3.5-5 is relatively small.

The Raman absorption of the system was also investigated, with beams approximately six times larger to address more of the ensemble. This was achieved with the same run that was used to investigate the EIT feature on resonance (see Tab. 11.3(b)(i)) but now with both the coupling and probe field centre frequencies moved off-resonance by the same amount (i.e., the one-photon detuning Δ). Fig. 11.9 shows the results from this investigation, with (a) showing the effect of changing Δ and (ii) showing the effect of changing the coupling field power. The first of these figures shows the expected behaviour of the Raman line, with the feature moving from enhanced transmission through to absorption as $|\Delta|$ is increased. For Fig. 11.9(b) however, it is normally expected that as the coupling power is increased the amount of absorption will increase. This was not the case here, with the amount of absorption staying constant at approximately 15% and the width slightly increasing as coupling power increased but being roughly 2 MHz. This low absorption again reflected a low optical depth, while the relatively large linewidth was most probably a combination of some residual magnetic field, and still being relatively close to resonance. In theory the latter point could be addressed by moving further off resonance (using a combination of the three AOMs and probe EOM present in the set-up) and increasing the power in the coupling field (remembering that $\text{OD}_{3l} \propto (\Omega_c/\Delta)^2$). However, the maximum power available for the coupling field in this set-up was 2 mW and with this it was impossible to see the Raman line further from resonance. The size of the coupling beam relative to the probe may have contributed as well, with a larger beam size potentially leading to improved absorption.

Time (s)	Action
42.00	Dipole trap shutter open
42.55	Repump amplitude off
42.5519	Upper MOT trapping beam amplitude off
42.553	Trigger for probe pulse
42.554	Repump amplitude on
42.5555	Repump amplitude off
42.565	Imaging shutter open
42.566	Push/Imaging beam amplitude on CCD Camera trigger
42.5661	Push/Imaging beam amplitude off
42.5681	Dipole trap shutter closed
43.0682	Push/Imaging shutter open
43.0692	Push/Imaging beam amplitude on CCD Camera trigger
43.0693	Push/Imaging shutter closed

Table 11.4: Light Storage Run for Expt.1d (Dipole Trap). The remainder of the run (following from Tab. 11.1) for the EIT light storage protocol attempted for Expt. 1d.

Finally, potentially the probe power could have been saturating the system and one option would be to reduce the probe power and see if the percentage of probe absorbed increased.

11.2.7 Interface and Detection Set-up 1d - Dipole Trap, $F = 1$

One final option to try to improve the optical depth of the Expt.1 set-up was to include a dipole trap beam through the centre of the MOT to attempt to increase the atomic density in this region. The laser used for the dipole trap was a SPI RedPower fibre laser with 20 W at 1070 nm focused down to a waist of approximately 50-100 μm in the centre of the trap. To ensure maximum overlap of the probe and dipole trap beams, they were aligned so as to be counter-propagating through the MOT. This was achieved through a mixture of orthogonal polarisations for the two beams as well as the use of a dichroic mirror that, in this case, reflected light at 1064 nm and transmitted light at 795 nm with nearly 100% efficiency. This set-up is shown in Fig. 11.10(a).

In this run, as described in Tab. 11.4, atoms were loaded from the MOT into the dipole trap simply by passing the dipole trap through the centre of the MOT. The red-detuned dipole trap would then attract atoms towards the centre of the trap (see Sec. 9.2.3). Due to the saturation of the imaging detector, the dipole trap could not be imaged this way. Instead, to be able to see the dipole trap, the MOT had to be turned off, allowing the atoms that were not in the dipole trap to fall away due to gravity. An image of the dipole trap achieved using this method is shown in Fig. 11.10(b). The reason the dipole trap did not appear to be straight into/out of the screen was that the imaging system was on a slight angle with respect to the dipole trap.

To determine whether this rather simplistic dipole trap loading method increased the optical depth of the system, the absorption of the atomic ensemble was investigated with, and without, the dipole trap on. The size of the probe pulse used to test the absorption

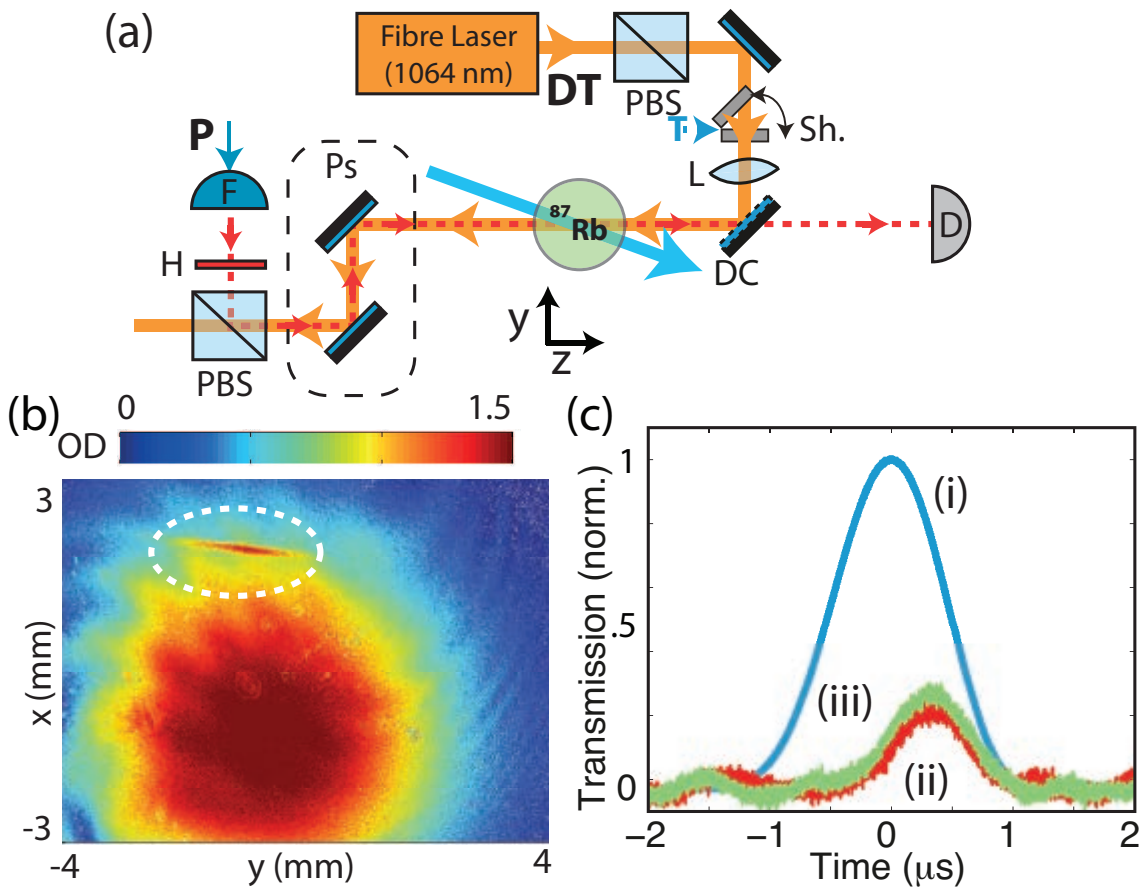


Figure 11.10: Interface and Detection Set-Up 1d - Dipole Trap. (a) Schematic diagram of set-up: P - probe field (dashed red line); I - imaging field (blue line); DT - 1064 nm 15 W dipole trap (orange line); F - fibre out-coupler; H - half wave-plate; PBS - polarising beam-splitter; Ps - periscope; L - lens; Sh. - shutter; T - digital triggers; DC - dichroic mirror; D - detector; ^{87}Rb - trapped ^{87}Rb atoms (MOT coils and trapping fields not shown). (b) Absorption image of dipole trap (circled). Colour bar shows OD scale. (c)(i) Probe pulse with maximum peak height of $4 \mu\text{W}$ (blue), (ii) absorption of the pulse without the dipole trap (red) and (iii) with the dipole trap (green).

was chosen to be slightly greater than the pulse area that could be totally absorbed by the system. Therefore if the dipole trap increased absorption, the amount of light leaking through to the detector would be reduced. As can be seen from Fig. 11.10, this was not the case, with the dipole trap making no visible difference to the absorption (here taken with a focusing lens in place). This is not surprising as this loading method from MOT to dipole trap was not very efficient. It would have been much better to mode-match the ensemble to the dipole trap first, but this would have made the process much more complicated and time was limited.

11.2.8 Experiment 1 Discussion

As stated at the start of this chapter, the aim of this work was not to produce the best possible EIT using cold atoms, but to identify experimental issues and tricks to help improve the system for envisaged future experiments. Therefore, though we were unable to demonstrate light storage using EIT, we cannot classify this work as a failure as many lessons were learnt.

The main issue that became apparent from this experiment was the difficulties presented when working with a finite number of atoms, as opposed to previous work carried out in warm vapour cells where the atom number appeared virtually infinite. The lack of optical depth was an issue raised multiple times in this section as the most likely reason for the short pulse delay times, as well as the small amount of Raman absorption. Developing a system with a high optical depth is therefore of critical importance for future experimental work, potentially using some of the techniques discussed in Sec. 9.4.2. However, a good first move would be to use a less complicated MOT set-up.

Apart from low optical depth, another issue was detecting the small amounts of probe light needed to avoid saturating the system (i.e., $< 1 \mu\text{W}$ when focused). This has implications for the detection set-up and the signal-to-noise ratio of the system. In the above experiment we used a single photo-detector with high gain and an amplifier ($\times 10$) before the signal was large enough to be viewed on an oscilloscope. The main issue with using high gain is that bandwidth is sacrificed. In this case, the detection bandwidth was on the same order as the signals being viewed. One option to improve this would be to move from a photo-detector to an avalanche photo-diode (APD), though great care would need to be taken to avoid saturating the APD with coupling field leakage. Another way to improve the detector resolution would be to move to heterodyne or homodyne detection. This method would require a local oscillator (LO) field with much greater power than the probe.

Another signal-to-noise issue that arose during the course of this experiment was to do with the SAS locking mechanism used. The SAS set-up made use of an internal 250 kHz modulation applied to the ECDL current. However, applying current modulation will cause amplitude and frequency fluctuations of the light that is sent to the main experiment, especially when it is sent through a cavity, as was the case for the probe field here. This is the main reason why the traces included in the above figures are so noisy, even after multiple averaging. It is therefore imperative for any future system design to avoid current modulation for the SAS lock. Zeeman modulation, for example, would be a much better alternative. The ring cavity shown in Fig. 11.1 was also locked using the laser current modulation and so, an additional EOM or another signal sent to the fibre-coupled EOM, would be needed to produce an error signal here.

Signal-to-noise could also be improved by further averaging. As can be seen from many of the figures in this section, most traces had four or more averages. The reason more averages weren't included in these traces was the approximately one minute run-time. A system can change a great deal in a minute. For instance, cavities and laser can come unlocked and powers through fibres can decrease due to polarisation rotation caused by minor fluctuations in room temperature. These can have large effects on the size of the MOT and also coupling field power. Shorter experimental runs are therefore desirable and, again, a less complicated system would be the good first step.

11.3 Experiment 2 - Caesium MOT

We will now discuss work carried out in the Laboratoire Kastler Brossel, at the Université Pierre et Marie Curie in Paris under the supervision of Dr. Julien Laurat and Prof. Elisabeth Giacobino. The ultimate goal of this work was to produce and store an entangled pair of photons in different parts of a caesium MOT (i.e., different “memories” as the speed of the atoms would not allow those from the different halves to interact) using EIT. More information on this, as well as a more in-depth description of the experiment can be found in Refs. [410, 411].

Though many elements, including the element, differ from those proposed for the cold atom Λ -GEM experiment, there are still techniques used here that could be applied to our future work.

11.3.1 Experiment 2 Set-Up

Here we will ignore the squeezers and homodyne detection set-ups, which were part of the experiment, and focus solely on the elements that were directly related to storage of weak coherent pulses using EIT. A basic schematic of the important experimental components is shown in Fig. 11.11(a). A Spectra-Physics Matisse titanium:sapphire (Ti:Saph) laser and an external cavity diode laser (ECDL) were phase locked 9 GHz apart (i.e., 200 MHz off the ground-state splitting of caesium, see App. A). This was achieved by taking a small amount of light from both lasers, beating them on a detector and mixing down at 8.6 GHz. This 400 MHz signal was then used to lock the phase of the two lasers with the help of a 100 MHz reference signal.

Unlike the first experiment presented in this chapter, the Ti:Saph laser was locked on resonance to the D2 $F = 4 \rightarrow F' = 4$ transition using a SAS set-up (not pictured) without an AOM to shift its frequency. As discussed in Sec. 11.2.1, this meant that, if we wanted to pulse the laser with an AOM, we would no longer be on resonance. To overcome this issue, two AOMs were placed in the Ti:Saph beam path: one to shift the frequency down, and the second to shift it back up to resonance. The ECDL path contained only one AOM, again for pulsing, which was driven at 200 MHz to bring the coupling and probe fields exactly onto two-photon resonance. The frequency of the ECDL could be altered by changing the locking point of the phase lock.

Also unlike previous experiments, neutral density (ND) filters were used as the primary method for controlling the powers in the different fields, as opposed to PBSs. This was because, with very small amounts of power transmitted through a PBS cube, small fluctuations in the polarisation of the fields would lead to a large percentage change in

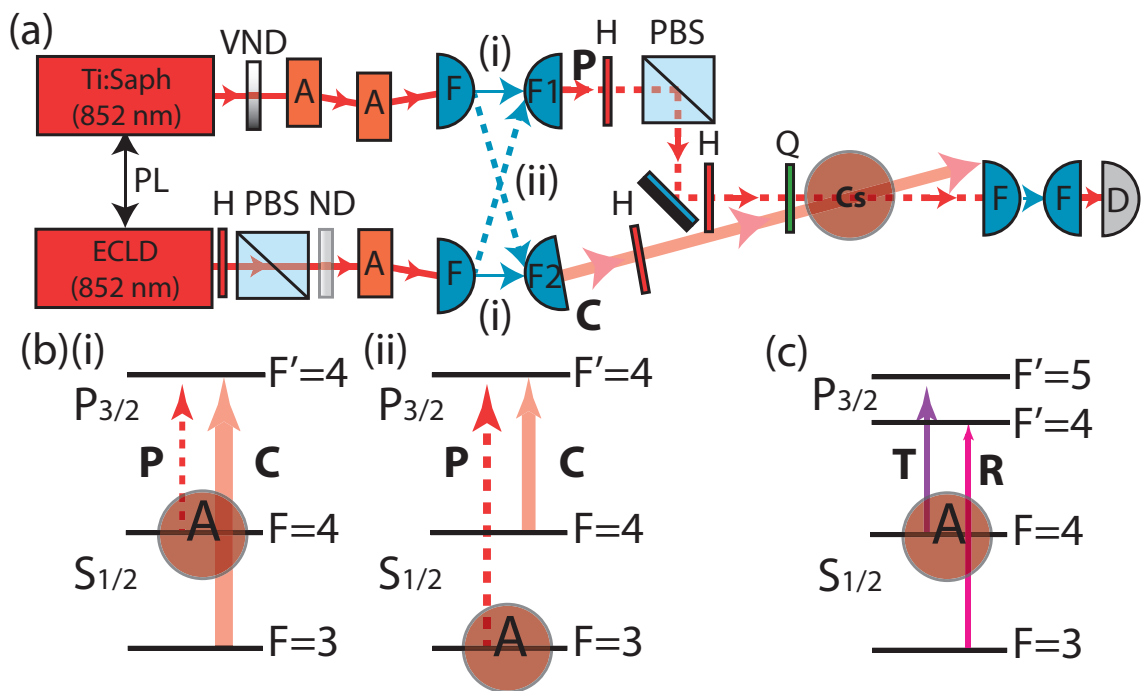


Figure 11.11: Cold Atom Cs Memory Set-Up, Expt. 2. (a) Schematic diagram of set-up: Ti:Saph - titanium:sapphire laser at 852 nm; ECDL - external cavity diode laser also at 852 nm; PL - phase lock between lasers (components not shown); H - half wave-plate; PBS - polarising beam-splitter; VND - variable neutral density (ND) filter; ND - fixed ND filter; A - acousto-optic modulator (drivers not shown); F - fibre couplers; P - probe field; C - coupling field; Q - quarter wave-plate; D - detector (either APD or high-gain detector); Cs - trapped Cs atoms (MOT coils and trapping fields not shown). (i) and (ii) refer to the two different experimental fibre configurations (solid and dashed lines respectively), corresponding to energy levels shown in (b)(i) and (ii). (c) Energy levels for trapping (T) and repump (R) fields with A indicating the location of the atomic population.

output power. The ND filters here were placed before fibre-coupling to ensure that the beam profiles through the MOT were not altered.

The trapping fields for the MOT were provided from two extra ECDLs (not pictured) and a Toptica tapered amplifier, with the trapping field containing 75 mW of power 10 MHz red-detuned from the D2 $F = 4 \rightarrow F' = 5$ transition and the repump field containing 4 mW of power on resonance with the D2 $F = 3 \rightarrow F' = 4$ transition. These transitions are shown in Fig. 11.11(c). The magnetic field was produced by a pair of anti-Helmholtz coils with a radius of 10 cm and 120 turns to create a gradient of 7.5 G cm^{-1} . The MOT was created in a glass cell at a pressure of around 10^{-9} Torr, with the Cs being supplied by a set of Cs-filled wires, as opposed to a dispenser. The MOT was loaded for 18 ms before each experimental run, with the runs repeating every $25 \mu\text{s}$. Approximately 10^9 atoms were trapped using this configuration.

The timing of the experiment was controlled via an FPGA-based LabVIEW code developed by S. Burks [410] that, among other functions, had the ability to switch the AOM driving frequencies on or off to remove the trapping fields before the memory sequence, as well as create square pulses for the probe and coupling fields via TTL signals.

The MOT was surrounded by a set of compensation coils. Unlike those in the previous experiment, these compensation coils were placed very close to the MOT and the relative currents in each pair of coils could also be tuned (i.e., $x_1 - x_2$, etc.) to cancel external gradient fields. Magnetic field compensation will be discussed more in Sec. 11.3.2.

As in Sec. 11.2.6, the coupling and probe fields were spatially separated, with an angle of roughly 2° between them. The two fields were oppositely circularly polarised, with the coupling beam having a waist of $200 \mu\text{m}$ and the probe having a waist of $50 \mu\text{m}$. A fibre after the MOT was used for greater suppression of the coupling field, with the output of the fibre sent to either a high gain photodetector constructed in-house, or a Perkin Elmer avalanche photo-diode (APD).

The fibre-coupled outputs of the Ti:Saph and ECDL could be sent to either (i) F1 and F2 respectively, or (ii) the other way around, depending on the experiment being performed. In the latter case, with the probe being derived from the ECDL and resonant with the D2 $F = 3 \rightarrow F' = 4$ transition, a $900 \mu\text{s}$ pulse of trapping light was used to pump the atoms into the $F = 3$ level. Other fields could also be sent to either of these fibre out-couplers. In all runs the probe and coupling fields were not sent into the MOT until 3 ms after the magnetic fields had been turned off.

11.3.2 Diagnostic Tools

In this section we investigate two diagnostic tools that, in addition to the absorption imaging presented in Sec. 11.2.3, could be used for improving a MOTs performance in terms of OD and magnetic field compensation.

Optical Depth Measure

We have already discussed one method of measuring OD in this chapter using an imaging beam to probe the entire MOT. That method had the advantage of also providing information on the shape and physical size of the MOT. However, in the end we are not interested in the peak OD, but the integrated OD over our probe pulse. To provide a better measure of this integrated OD, one can send a near-resonant probe pulse through the MOT and measure its attenuation. This attenuation can then be scaled back to give the integrated resonant OD for the probe transition using Eq.s 11.1 and 11.2.




Figure 11.12: Optical Depth Measure, Expt. 2. Probe pulse transmission, (i) without (blue), and (ii) with (red) a MOT.

An example of an OD measurement using this technique is shown in Fig. 11.12, with (i) the reference pulse and (ii) the attenuated probe. The probe field in this case was derived from the locking laser, though using a separate AOM to allow for pulsing, and was 10 MHz detuned from the D2 $F = 4 \rightarrow F' = 5$ transition. This signal was then sent to fibre out-coupler F1. Here we found $I_t/I_o = 0.25$, giving an on-resonance OD of approximately 20 for the MOT (remembering $\gamma_{cs} = 5.2$ MHz).

Microwave Pulse Magnetic Field Compensation

One problem encountered in the previous experiment was residual magnetic fields broadening the EIT absorption line. This is a problem as external magnetic fields, either dc or gradients, will lead to shorter coherence times for the system (see Sec. 5.3.6).

In this experiment, the optimisation of the magnetic field compensation was achieved with the use of a microwave signal generator and the more complex compensation coil design discussed in Sec. 11.3.1. Firstly, an antenna was connected to the function generator and placed next to the MOT. A 900 μs microwave pulse was then sent into the MOT to create a π -pulse. One microsecond later, a 30 μs probe pulse on the D2 $F = 3 \rightarrow F' = 4$ transition was sent through the MOT, followed by a 200 μs repump pulse and finally a 30 μs reference pulse. If the microwave pulse was on resonance with a $|F = 4, m_F = n\rangle \rightarrow |F = 3, m_F = m\rangle$ transition then atoms would move from the $F = 4$ to $F = 3$ level and attenuate the probe pulse.

To investigate the splitting of the m_F states, and therefore the background magnetic field, the frequency of the microwave pulse was scanned and the OD of the probe pulse measured using Eq. 11.1 for all frequencies. This was all carried out automatically using the LabVIEW code mentioned in Sec. 11.3.1, which was capable of performing such a sweep quickly (i.e., less than one minute for low resolution) while recording I_t/I_o . Performing many such sweeps, while manually altering the currents in the six compensation coils, allowed for quick compensation of external fields.

Some example traces are shown in Fig. 11.13, with greater compensation moving from (i) \rightarrow (vi). The separation of peaks indicated the magnitude of the dc magnetic field to be compensated, and the width of the peaks indicated the magnitude of the magnetic field

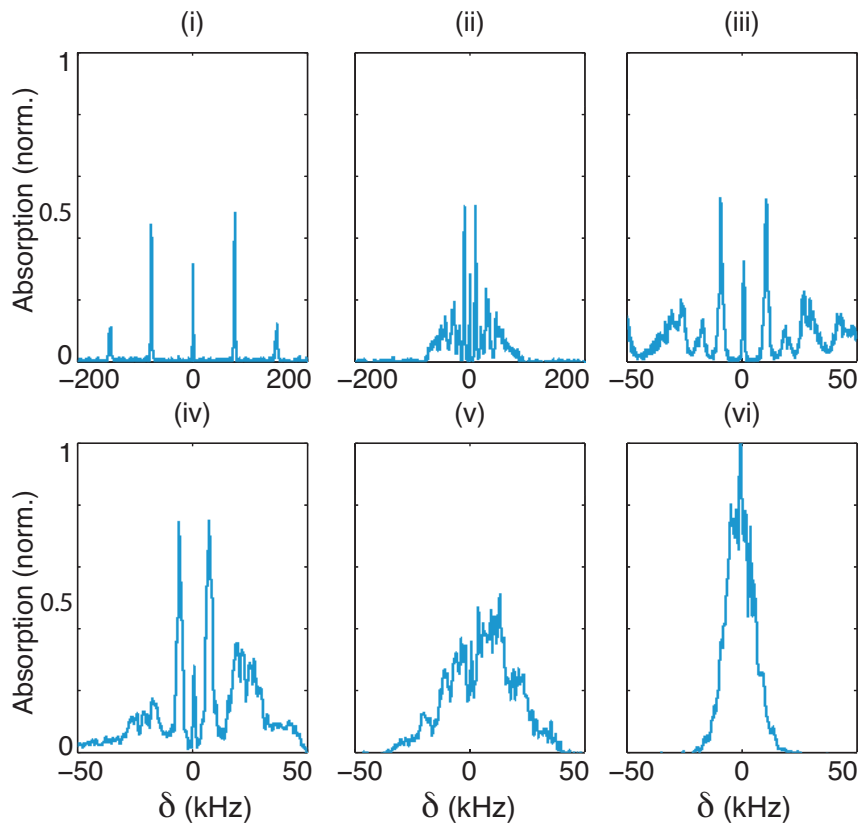


Figure 11.13: Magnetic Field Compensation, Expt. 2. Absorption of a probe pulse through the MOT, normalised to maximum absorption, as a function of microwave pulse frequency. (i)→(vi) denote the optimisation of the currents in the compensation coils.




Figure 11.14: EIT Delay and Storage, Expt. 2. Transmission of a probe beam through a Cs MOT in the presence of a coupling field. (a) Coupling field on for entire run. (b) Coupling field turned off from 0.5-1.5 μs . (i) Traces without the MOT present (blue), and (b) with MOT present (red). Approximately one photon per pulse in the probe field (detected using an APD) averaged over many traces, and approximately 100 μW in the coupling field.

gradient field to be compensated. As can be seen here, the final field after compensation gave a single absorption peak with a FWHM of approximately 10 kHz. This indicates a cancellation of magnetic fields down to 5 mG.

11.3.3 EIT Experiments

In this section we will look at two EIT-based experiments: the first based on storage of light using EIT; and then one that revolves around the transparency window and its origins.

Pulse Delay and Storage

With the much higher OD for the Cs MOT, and much smaller intensities that could be detected using the APD, we could finally observe EIT delay and storage, using pulses down to the single-photon level. Firstly, using set-up (i) from Fig. 11.11(a), and corresponding level structure shown in Fig. 11.11(b)(i), the coupling field was turned on approximately 3 ms after the MOT fields were turned off. This was followed 2 μs later by a 0.5 μs probe pulse. The polarisation of the probe and control were altered to try to maximise transparency using this method, with relatively high power in the coupling field. It is interesting to note here that it was found that the best polarisation were either same-circular (i.e., either both right, or left, circularly polarised) or opposite-linear, both giving maximum transmission of around 75%. While opposite-circular and same-linear gave less than 30%.

Probe pulse delay was seen by turning down the coupling field power, as shown in Fig. 11.14(a). By turning off the coupling field at the end of probe pulse, storage of this delayed probe was observed (Fig. 11.14(b)). With this set-up a maximum efficiency of 15% was achieved after storage for 1 μs . This agrees with the theoretical maximum efficiency for the measured OD of approximately 20 [412]. To attempt optimisation, we

wrote additional LabVIEW code that integrated the area under the input and echo to calculate the efficiency in real-time.

Experimental Investigation of the Transition Between Autler-Townes Splitting and EIT

With the Cs MOT we examined the transmission window for the probe field for a variety of different coupling field powers. The reason we were interested in this was to investigate the transition between EIT and Autler-Townes splitting (ATS) [221, 222]. Both effects produce a transmission window in a three-level system for a weak probe field, due to the presence of an additional coupling field. EIT arises due to Fano interference [220], where two modes that are almost resonant decay through a common channel with a weak coupling field (see Sec. 3.3). The EIT transmission window then becomes the difference of two Lorentzians, one much narrower than the other. On the other hand, ATS is the result of the splitting of an excited state in the dressed-state picture (see Sec. 2.5.2), due to a relatively strong coupling field. The ATS transmission window then becomes the sum of two Lorentzians with approximately the same widths.

Though both produce similar effects, it is important to be able to differentiate between the two as a transmission window alone does not guarantee coherent population trapping (the corresponding dynamics of atoms due to Fano interference [414]) as EIT does. The EIT transmission window is much sharper than the ATS window and therefore more sensitive to variations in the system, important for quantum metrology applications (as outlined in Sec. 1.3.4). Also, as ATS is not a coherent effect, storage in a window that is partly EIT and partly ATS could lead to increased decoherence. However, it has recently been shown theoretically that an ATS storage medium could be better for squeezed state storage than EIT [415]. Characterising the transition is also topical as in the past few years there have been experiments demonstrating transparency windows in a variety of novel systems such as in an “artificial atom” created in a superconducting quantum system [416], as well as in an opto-mechanical system using phonon-phonon interactions [417], and even using a vacuum coupling field with cold atoms in a cavity [418].

This had led to an interest from theorists to develop tests for whether a system can best be described by EIT or ATS [419, 420]. Here we used the method of P. M. Anisimov et al. [421] who, in 2011, introduced a test for differentiating between EIT and ATS using Akaike’s information criterion [422]. This is achieved by fitting the transmission window using both an EIT and an ATS model, and then using the NonLinearModelFit function in Mathematica to determine the Akaike weight W_A of each fit. This weight allows for a sensitive test of whether the system is in an EIT regime or not. The models for the two different systems are given by

$$A_{EIT} = \frac{C_+^2}{\gamma_+^2 + \delta^2} - \frac{C_-^2}{\gamma_-^2 + \delta^2} \quad (11.4)$$

$$A_{ATS} = C^2 \left(\frac{1}{\gamma^2 + (\delta - \delta_o)^2} + \frac{1}{\gamma^2 + (\delta + \delta_o)^2} \right), \quad (11.5)$$

where the C s represent Lorentzian amplitudes, with the γ s denoting the corresponding widths, and δ_o being the shift of the ATS peaks from two-photon resonance $\delta = 0$. For more information see Ref. [421].

To view the transmission window we needed to be able to scan the two-photon de-




Figure 11.15: EIT to Autler-Townes Transition, Expt. 2. (a) Enhanced transmission windows for a variety of coupling field Rabi frequencies Ω_c as a function of two-photon detuning δ , with absorption being plotted in arbitrary units. (b) Single transmission windows for (i) $\Omega_c = 0.2\gamma$ and (ii) $\Omega_c = 2.3\gamma$ as a function of δ . Points represent experimental data, the red line represents EIT fit to data and the green line represents Autler-Townes fit. Adapted from Ref. [413].

tuning. However, with the set-up described in Fig. 11.11(a)(i) it was not possible to scan the Ti:Saph laser frequency, as a result of the double AOM configuration. Instead, the ECDL was used for the probe, and the Ti:Saph used as the coupling field, as shown in Fig. 11.11(a)(ii). Even with this change it was not possible to automatically scan the frequency of the ECDL using its AOM, as the single-pass nature of this AOM meant that changing the frequency would significantly alter the probe power going through the MOT. Instead, δ was scanned by hand by changing the frequency of the phase lock. The absorption measurements were made for individual pulses of approximately 30 nW, as in Fig. 11.12, with the ECDL frequency altered between pulses. In this arrangement opposite circular polarisations were used for the probe and coupling fields.

To characterise the transition from EIT to ATS we investigated the transmission window for a variety of coupling field powers (100 nW \rightarrow 200 μ W) with a MOT OD of three to avoid any bandwidth issues with the detector. The absorption profiles obtained are plotted in Fig. 11.15(a). As can be seen, as the power in the coupling field P_c (and Rabi frequency $\Omega_c \propto \sqrt{P_c}$) increased, the width of the transmission window increased. Figure 11.15(b) shows two of these transmission windows for (i) low and (ii) high coupling field power, as well as fits to the data using modified versions of Eq.s 11.4 and 11.5. As can be seen, for a low coupling field Rabi frequency of $\Omega_c = 0.2\gamma$, the EIT fit is much better than the ATS fit, with Akaike weights of $W_A = 0.85$ and 0.15 respectively. For the much higher coupling field Rabi frequency of $\Omega_c = 2.3\gamma$ the ATS fit was better than the EIT fit, with $W_A = 0.65$ and 0.35 respectively.

Implementing the same procedure as above for the full range of coupling powers produced a curve similar to that in Fig. 3 of Ref. [421]. The experimental transition point between EIT and ATS (i.e., where $W_A = 0.5$ for both) was found to be at $\Omega_c = 1.23\gamma$. This did not agree with the simple three-level theoretical description. However, using a multi-level approach (i.e., six separate Λ systems for the Cs atoms) [423, 424], and assuming an inhomogeneous broadening of 1.3 MHz produced a theoretical model that agreed with the data. For more information see Ref. [413].

11.3.4 Experiment 2 Discussion

A much larger OD atomic ensemble, far shorter loading times, and more sensitive detection meant that this system far out-performed the previous experiment in terms of EIT storage and signal-to-noise measurements. It also, more importantly, provided more lessons for future cold atom experiment development.

Both the pulsed OD and microwave pulse magnetic field measurement methods are very useful tools for investigating and improving a cold atomic memory's operation. The pulsed OD method could prove useful for fine alignment of the probe through an ensemble, while having a quick and accurate method for cancelling the background magnetic field could help to reduce decoherence in the memory. However, the complicated nature of the microwave pulse measurements and the additional components required (namely an additional high frequency function generation, additional LabVIEW hardware and coding, and an antenna), coupled with the fact that GEM is less sensitive to small fluctuations in the magnetic field due to the large dc offset magnetic field required, means that implementing this type of scheme is not of immediate concern.

Another useful technique learnt from this experiment was the use of ND filters, rather than relying on PBS cubes, to reduce power in the different fields, reducing oscillations in

power caused by fluctuations in polarisation. The phase lock provided far greater power than was achievable with the coupling field in the previous experiment. However, the phase of the two lasers did not remain locked for long periods (i.e., greater than ten minutes) and required constant attention. Another constraint of this experiment was the inability to shape the probe pulse, as there were no arbitrary waveform generators used. The square pulse, produced by simply turning on and off the AOM driving fields, had a larger bandwidth than a Gaussian pulse of roughly the same temporal length, reducing the maximum storage efficiency possible.

Investigating the transition from EIT to ATS by looking at the transmission windows is another illustration of why cold atomic ensembles can be so useful for quantum memories: the cancellation of surrounding magnetic fields to 5 mG and low temperature of the atoms dramatically reduced the inhomogeneous broadening of the ensemble. A larger broadening would have led to a smaller saturation difference in Akaike weights for the two models. However, not being able to easily scan both the probe and coupling frequencies proved an issue when looking at the transparency windows. The addition of a double-pass AOM in the ECDL line, with a triggerable function generator providing the driving signal would allow for a scan in 1 ms (as was the case with Expt. 1), rather than a process taking tens of minutes. This would have allowed for more values of Ω_c to be sampled, and less noise per trace.

11.4 Conclusions

The first, three-tier ^{87}Rb , MOT was not designed for a quantum memory experiment due to the long loading time required for the upper MOT and the low OD of five. It is therefore not overly surprising that it did not produce large EIT delays that, in turn, ruled out the possibility of storing light. However, lessons learnt about the techniques such as Zeeman modulation, absorption imaging, as well as methods (to either implement or avoid) for improving OD and signal-to-noise, have shaped the proposed development of the cold atom Λ -GEM experiment presented in the next chapter.

The second, Cs, MOT was designed with quantum memory applications in mind and therefore not unsurprisingly gave much better results. This included a probe OD of 20, and storage of information with up to 15% efficiency for a 1 μs storage time. With this MOT we were also able to investigate the transition from EIT to Autler-Townes splitting, finding a transition point at $\Omega_c = 1.23\gamma$. The behaviour matched the theory, as long as all hyperfine levels of the caesium atom were considered and an inhomogeneous broadening of 1.3 MHz was assumed. Useful techniques gleaned from this experiment included the pulsed probe OD measurement, automatic microwave pulse measurements of the magnetic field and the use of ND filters to lower power fluctuations in the fields, which again helped in the development of the cold atom GEM experiment.

Cold Atom Gradient Echo Memory

Kaylee: “Shiny” - FireFly (repeated line), Fox, 2002-2003, TV Series.

12.1 Introduction

In this chapter we present a cold atom realisation of the Λ -GEM scheme, motivated by and building upon the work presented in Ch.s 9-11. To achieve this we first had to develop an ultra-high optical depth atomic source to allow for high efficiency recall. Three different experimental configurations were attempted to optimise the memory. The chapter is split into three main sections: 12.3; 12.4; and 12.5, each describing one of these configurations. Preceding this, we present the basic set-up common to all configurations in Sec. 12.2.

Most of the work presented in this chapter has been published in the journal article:

- B. M. Sparkes, J. Bernu, M. Hosseini, J. Geng, Q. Glorieux, P. A. Altin, P. K. Lam, N. P. Robins, and B. C. Buchler, *Gradient echo memory in an ultra-high optical depth cold atomic ensemble*, arXiv:1211.7171v3 (2012).

12.2 Common Experimental Set-Up

The cold atom GEM experiment spanned two tables: one containing all lasers, cavities and AOMs; and one containing the vacuum system, cold atomic source and detection. The two tables were linked via optical fibres.

12.2.1 Table 1

GEM Set-Up

The set-up of the coupling and probe fields was the same as those shown in Fig. 11.1 with the modifications listed below. Firstly, an M2 SolsTiS titanium:sapphire laser (Ti:Saph), with an external reference cavity and a 12 W Laser Quantum Finesse pump, was used rather than a diode laser to produce 3 W of power at 795 nm. The laser could either be placed near the D1 $F = 2 \rightarrow F' = 2$ transition, with the frequency being stabilised by the reference cavity, or locked to the transition using saturated absorption spectroscopy

(SAS). In the latter case, Zeeman locking [408] was used to create the modulation for the SAS-peak error signals, with a modulation frequency of 250 kHz, rather than modulating the laser directly. The error signals were sent to a modified version of the LabVIEW code described in Ch. 4, and then the output control signal was sent to the reference cavity of the Ti:Saph. Next, by adding a beam-splitter after the ring cavity, the probe set-up was duplicated to also create a local oscillator (LO) field. Finally, rotating neutral density filters were used to precisely control the power in the probe and coupling fields, as discussed in the previous chapter.

Trapping and Repump Set-Up

The MOT field set-up was, again, similar to the one discussed in the previous chapter. The trapping and repump fields for the MOT were created with two MOGLabs external cavity diode lasers (ECDLs) and a bare laser diode (also stabilised via a MOGLabs controller). The first ECDL was used to create the repump field. This was current-modulated to allow for locking to the D2 $F = 1 \rightarrow F' = 2$ 1,2 crossover using SAS, with 22 mW present after a single-pass acousto-optic modulator (AOM) and fibre-coupling. The second ECDL was not modulated, with Zeeman modulation used to lock to the D2 $F = 2 \rightarrow F' = 3$ 2,3 crossover using SAS (see Sec. 11.2.1). This laser was used to create two imaging fields from one double-pass AOM, with a total power of only a few milliwatts after fibre-coupling, as well as to injection-lock the laser diode. This laser diode was passed through a M2K tapered amplifier to create the trapping field, with 130 mW present after a double-pass AOM and fibre-coupling. All AOMs in the MOT field set-up were controlled via Minicircuits voltage controlled oscillators and attenuators, which were in turn controlled via a modified version of the LabVIEW described in Sec. 11.2.2.

Figure 12.1(a) shows the frequencies of all fields present in the experiment. For more information on the set-up of the first optical table see Fig. 12.14 at the end of the chapter.

12.2.2 Table 2

Vacuum System

The centre-piece of the second table was the vacuum system used to create the cold atom cloud. This system consisted of two six-inch-diameter, ten-inch-long nipples placed either side of a double-sided anti-reflection (AR)-coated ultra-high vacuum glass cell from JapanCell¹ with dimensions of $20 \times 20 \times 76$ mm, with the long axis in the y -direction. This is shown in Fig. 12.1(b). Optical access through the vacuum system (i.e., down the y -axis) was provided by a roughly two inch double-sided AR-coated window on the other side of the vacuum system. To achieve as high a vacuum as possible, the system was pumped with a roughing pump and baked to approximately 140°C to remove as many impurities as possible. A 70 L s^{-1} Gamma Vacuum ion pump was constantly operated at one end of the vacuum system and a Duniway titanium sublimation pump on the other side was used to reduce the background pressure to approximately $5 \cdot 10^{-10}$ Torr before the vacuum

¹Note: Care should be taken when baking with these cells, as the glass sides are stuck together using an epoxy that has not been cured and, therefore, it must be baked at atmosphere to at least 160° before being baked under vacuum.

systems initial use. The rubidium was supplied by a natural mixture Rb Alvatec dispenser.

Compensation Coils

Compensation coils and a non-magnetic table were also used to create a null magnetic field for the atoms. This was important for both the atomic cloud creation, and memory, processes. The compensation field was created by three sets of Helmholtz compensation coils in the x -, y - and z -directions. Due to the dimensions of the table, the radii of these coils were 700 mm in the x - and z -directions, and 1 m in the y -direction. All coils had 18 turns.

12.3 Configuration 1 - MOT Fields at 45° to GEM Fields

To create as high an OD as possible we optimised both the static loading of the MOT through geometry, a spatial dark spot and optical de-pumping, followed by a compression phase using a temporal dark spot, all as described in Ch. 9. After describing the application of these techniques, we present a characterisation of the system using near-resonance absorption imaging. We then present the set-up of the GEM fields and results achieved with this MOT set-up.

12.3.1 MOT 1 Set-Up, Optimisation and Characterisation

Our first ^{87}Rb MOT was in a three-beam retro-reflection configuration. The trapping and cooling laser was initially red-detuned by 25.5 MHz from the D2 $F = 2 \rightarrow F' = 3$ transition for the loading of the MOT. The repumping field was on resonance with the D2 $F = 1 \rightarrow F' = 2$ transition. These are shown in Fig. 12.1(a). Figure 12.1(b) shows the MOT beam configuration, with trapping and repump beams being combined on a polarising beam-splitter (PBS) and then split further using another PBS (neither pictured) to create three cooling beams with diameters of approximately 2.5 cm, all of which were then appropriately polarised and retro-reflected. Further detail is shown in Fig. 12.15.

The optimal shape for a cold atomic ensemble for use as a quantum memory is a cylinder along the direction of the memory beams to allow for maximum absorption of the probe. This was achieved here by using rectangular, rather than circular, quadrupole coils, as described in Ch. 9, with dimensions of 85×40 mm and 45 turns in each coil, placed against the glass cell. The shape of the MOT could then be altered by varying the current in the coils, as well as the relative intensities of the trapping fields.

For the loading phase, the MOT coils produced a gradient of 37.5 G cm^{-1} . The density in the trapped atomic state ($F = 2$ here) was limited by reabsorption of fluorescence photons within the MOT, leading to an effective outwards radiation pressure [356]. By placing a dark spot of approximately 5 mm in diameter in the repump beam, atoms at the centre of the MOT were quickly pumped into the lower ground state ($F = 1$) and become immune to this unwanted effect, allowing for a higher density of atoms in the centre of the trap. We collect over 10^9 atoms in this configuration.

In the compression stage of the ensemble preparation we used a temporal dark spot [370, 425, 399] to transiently increase the density of the sample by simultaneously increasing the trapping laser detuning to -56 MHz, decreasing the repump intensity to approximately 20%, and reducing the magnetic field gradient by 13%. Detuning the trapping




Figure 12.1: Cold GEM Set-Up 1. (a) Energy level diagram for all fields present: C - coupling; P - probe; LO - local oscillator; I - Imaging; T - trapping; R - repump. Also shown are one- and two-photon detunings for probe and coupling fields (Δ and δ respectively). (b) 3D representation of MOT configuration, with: TR - trapping and repump fields; MC - MOT coils (green). (c) 2D schematic diagram of experiment, with: ^{87}Rb - atomic ensemble; BS - 50:50 beam-splitter; L - lens; D - photo-diode detector; I1(2) - imaging beam 1(2); CCD - charged-coupled device camera; Ir - iris; GC - GEM coils (H - Helmholtz, AH - anti-Helmholtz); Ps - position for inserting avalanche photo-diode. MOT coil configuration not shown, neither is vertical (x) MOT beam.




Figure 12.2: Imaging Frequency Calibration. Peak OD (normalised to maximum) as a function of frequency of the imaging beam AOM driver. Points show data averaged over five traces, the line shows a fit to data.

laser creates some of the conditions for polarisation-gradient cooling (PGC - see Sec. 9.2.2). The compression stage lasted 20 ms.

Finally, we transferred the atomic population into the $F = 1$ level by simply turning off the repump 1.4 ms before the trapping beam (see Fig. 12.1(a)). Once in the $F = 1$ level we attempted to optically pump atoms into either of the $m_F = \pm 1$ states using the method described in Ref. [403], except with lasers on the D2 transition rather than the D1. In this case we were only able to place up to 50% of atoms in the correct state and therefore did not use this in our final GEM realisation.

Absorption imaging, as described in the previous chapter, was used to optimise and characterise the MOT. The set-up and frequency of the imaging beams are shown in Fig. 12.1. We performed imaging both across (I1) and at a slight angle to (I2) the z -axis. To measure the detuning for the imaging fields accurately we lowered the atom number in our trap until the OD did not saturate on resonance and plotted out the resonance curve as a function of detuning to locate the line centre. This characterisation is shown in Fig. 12.2. Note that, as it is a double-pass AOM set-up, the detuning from resonance should be multiplied by two.

Images of the MOT from both front- and side-on without the dark spot are shown in Fig. 12.3. For these imaging runs, we used a 4.98 s load time before ramping the optical frequencies and coil currents as described above. All fields were then turned off and an image of the MOT was taken 1.5 ms later. As the imaging beam was on the closed D2 $F = 2 \rightarrow F' = 3$ transition, the atoms were not pumped to $F = 1$ for the imaging runs. A comparison image was taken 150 ms later to obtain I_o (see Eq. 11.2). The imaging beam was detuned by -20 MHz to avoid complete absorption and therefore saturation of the measured OD.

The MOT shown in Fig. 12.3 had $1.05 \cdot 10^9$ atoms with the width in the z -direction being nearly twice the x - y width. The reason for the elliptical front-on view was that the two rectangular coils did not give a symmetric magnetic field in the z -direction. This led to a peak OD of 110 on the D2 $F = 2 \rightarrow F' = 3$ transition. With the addition




Figure 12.3: MOT 1 Characterisation Images. (a) Side-on and (b) front-on images of the atomic ensemble for the optimised MOT. Colour bars show OD scales, σ values show standard deviation of Gaussian fits to ensemble.

of the dark spot, as mentioned in the previous section, the peak OD increased to 140. The temperature of this ensemble, determined by measuring the width of the expanding cloud 1.5 and 5 ms after the MOT fields were turned off, was approximately $200 \mu\text{K}$ in the x,y -plane, and $50 \mu\text{K}$ in the z -direction. The difference in temperatures can be attributed to the fact that the MOT was confined much more tightly in the x,y -plane than in the z -direction and, therefore, we would expect a higher temperature in this plane.

12.3.2 GEM Set-Up and Results with MOT 1

The set-up of the GEM fields relative to the MOT fields is shown in Figs 12.1(b) and (c), with more detail being shown in Fig. 12.15. The alignment of the probe beam to the atomic ensemble was crucial to achieving the highest on-resonance OD for the memory, and therefore the highest efficiency. The probe beam was focussed into the MOT to access the highest optical depths available in the atom cloud. Using lens L2 we formed a waist of $200 \mu\text{m}$ in the atom cloud. The probe was aligned to the MOT using a combination of the imaging technique, discussed in Sec. 11.2.3, and pulsed OD measurements, discussed in Sec. 11.3.2, with a detuning of -30 MHz for the probe. The latter measurements were taken using a high gain Thor Labs photo-detector, placed at position Ps. With the alignment optimised we measured a resonant probe OD of 100 without the dark spot, and 120 with the dark spot.

To combine the probe and coupling fields we used a non-polarising 50:50 beam-splitter (BS1 in Fig. 12.1). Unlike the probe, we wanted the coupling field to have a large diameter to cover the entire ensemble with minimal variation in intensity. We used L1 to make a telescope with L2 (the focusing lens for the probe) to have a collimated coupling field with a waist of approximately 5 mm. The reason we used a BS rather than a PBS was that it allowed us to vary the probe and coupling polarisations, to which GEM is very sensitive [252]. For quantum signals such as squeezed states, where loss of the probe before the MOT is important, the 50:50 beams-splitter could be replaced with a 95:5 beam-splitter, as only 1-2 mW were ever used for the coupling field and we had the potential for up to a watt of coupling power available on Table 2. After the MOT, we filtered the coupling field from the probe using an iris with a diameter of 1 mm. This removed most of the coupling field while allowing all the probe light through. The mode-selective heterodyne detection registered the remaining coupling field as a small dc offset.

As shown in Fig. 12.1(c), the GEM fields were created with four coils: two Helmholtz with seven turns each; and two anti-Helmholtz coils with 12 turns each. All coils had a radii of 35 mm and the two sets were placed 40 mm apart. The magnetic field offset was created by placing the same current in both Helmholtz coils. The output gradient was then created by putting more current into one of these coils and less into the other, and the anti-Helmholtz coils were turned on to create the input gradient in the opposite direction. This design was used so that currents would only be turned off to reverse the gradient, similar to the variable-pitch solenoid configuration used for the warm vapour cell experiments (see Ch. 6). This led to switch-off times on the order of microseconds, with only small residual oscillations.

To perform GEM with this set-up, we used the MOT parameters discussed in Sec. 12.3.1 with a 480 ms load time. After this time the fields were ramped down, as discussed above, and then all MOT fields were turned off. At this point the input gradient was turned on, using all four GEM coils. 500 μ s later, after allowing most of the eddy currents generated by the MOT coils to decay, the coupling field was turned on and then the probe was pulsed into the ensemble. Though the MOT could not fill completely in the 480 ms load time, by continually running the above sequence and having only a millisecond between turning off the MOT and turning it back on again, the atom number in the MOT saturated after only five cycles of the experiment.

We found a Gaussian pulse with a full-width-half-maximum (FWHM) of 10 μ s to be optimal for storage. This was a trade-off between decoherence effects for longer pulses, and the higher bandwidths (and therefore reduced ODs - see Eq. 3.54) required for shorter pulses. For this pulse length, a one-photon detuning of approximately -250 MHz, a coupling field power of roughly 1 mW, and a peak probe power of under 1 μ W, we were able to demonstrate storage with a maximum of 70% total efficiency. This was measured by both fitting modulated Gaussians to the input and echo, as shown in Fig. 12.4(a), and using the amplitude and width to determine the Gaussian's area, as well as squaring the data and integrating (using the Mathematica code presented in App. D). From the modulated heterodyne data we could also see that the echo had approximately the same frequency as the input pulse to within 50 kHz.

To investigate the coherence time, we delayed the gradient switching time to store the pulse for longer periods inside the memory while turning off the coupling field during the storage window. This data is shown in Fig. 12.4(b) for storage both 500 μ s, and 2.5 ms, after the MOT fields were turned off. Both showed a Gaussian decay with time constants of 208 μ s and 254 μ s respectively, with 0.4% efficiency for storage past 500 μ s seen for the latter. The much lower initial efficiency for the storage after 2.5 ms was due to the reduction in MOT OD caused by thermal expansion. The Gaussian nature of the decay indicated that the storage time was most probably diffusion-limited, however the difference between the two MOTs, which had the same temperature (and therefore same expansion rate), indicated that there were other effects, most probably eddy currents and associated magnetic field decoherence, present 500 μ s after the MOT fields were turned off, but had decayed by 2.5 ms.

The Gaussian nature of the decay meant that for short storage times the efficiency was relatively constant, and so multiple pulses should be able to be stored with the same efficiency. This is illustrated in Fig. 12.4(c), where two input pulses, separated by 30 μ s were stored and then retrieved. Due to the time-reversal nature of GEM, the first echo out (iii) comes from the second input pulse (ii) and so there was a large difference between the storage times of 40 μ s for input (ii) and 85 μ s for input (i). However, the efficiencies,

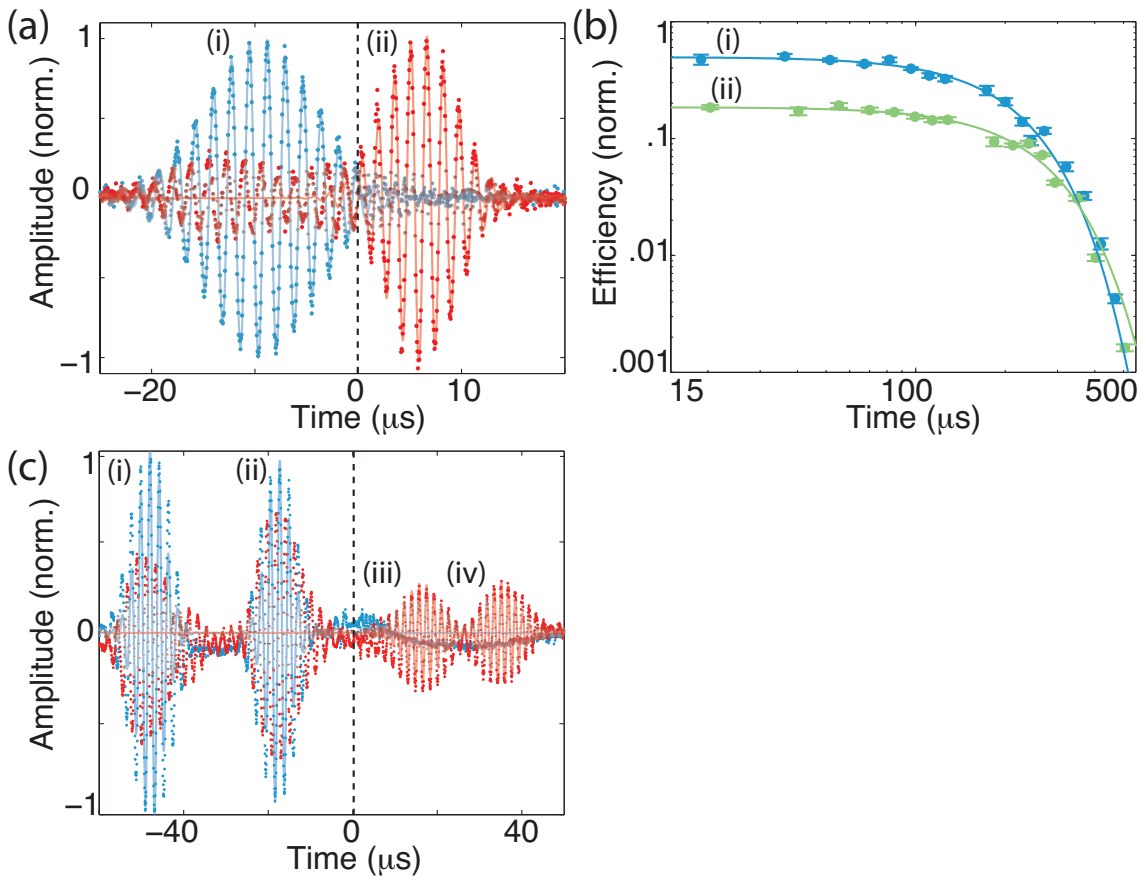


Figure 12.4: MOT 1 GEM Results. (a) High efficiency modulated heterodyne data: (i) input pulse (blue), and (ii) echo (red) with 70% total efficiency. Points correspond to single-shot data, lines correspond to modulated Gaussian fits. The dashed line indicates magnetic field switching at $t = 0$. (b) Log-log plot of memory efficiency as a function of storage time for the MOT shown in Fig. 12.3 (i) 500 μs after MOT turn-off (blue), and (ii) 2.5 ms after MOT turn-off (green). Points correspond to data averaged over 10 traces, error bars show standard deviation, and lines show Gaussian fits to the data. (c) Multi-pulse storage and retrieval: (i) input pulse 1 (blue), (ii) input pulse 2 (blue), (iii) echo 1 (red), (iv) echo 2 (red). Points correspond to modulated data, and lines correspond to fit to the data. The dashed line indicates magnetic field switching at $t = 0$.




Figure 12.5: Cold GEM Set-Up 2, with: ^{87}Rb - atomic ensemble; TR- trapping and repump fields; I1(2) - imaging beam 1(2); C - coupling field; P - probe field; LO - local oscillator field; BS - beam-splitter; L - lens; GC - GEM coils; PH - pin-hole; D - detector; CCD - charged-coupled device camera. MOT coil configuration not shown, neither is vertical (x) MOT beam.

though small, were approximately the same with 8% for echo (iii) and 5% for echo (iv). This loss will also be partly due to coupling-field-induced scattering (see Sec. 5.3.4), not present in the decay shown in Fig. 12.4(b).

12.4 Configuration 2 - MOT Fields Parallel and Perpendicular to GEM Fields

As can be seen from Fig. 12.4(a), there was significant leakage of the probe beam, indicating that if we could increase the OD of the probe we could improve efficiencies. To try to achieve this we decided to optimise the MOT and the set-up further.

12.4.1 MOT 2 Set-Up, Optimisation and Characterisation

To increase the overall number of atoms we could trap we needed a larger trapping volume. To achieve this the glass cell was replaced with a larger ($100 \times 50 \times 50$ mm) single-sided AR-coated cell from Precision Glass Blowing, including another natural mixture Rb dispenser, this time from SAES. To keep the intensity of the trapping beams high, and allow us to detune further from resonance during the compression stage, we replaced the tapered amplifier mount from the previous set-up with a custom-designed one to improve seed alignment, and therefore power in the MOT beams, with over 400 mW of trapping field power then available after the fibre.

We also altered the MOT geometry relative to the GEM fields, with our idea being to elongate the MOT further along the direction of probe propagation to increase the interaction volume. This was achieved by placing the GEM fields down the y -axis (as opposed to the z -axis), i.e., the long-axis for the vacuum cell. The basic set-up is shown in Fig. 12.5. The main changes are listed below.

Firstly, rather than using three pairs of retro-reflected beams, we decided to use two separate beams in the direction of the GEM fields (y -direction). This allowed us to have fine control of the length and position of the MOT in the y -direction by changing the




Figure 12.6: MOT 2 Characterisation Images. (a) Side-on and (b) front-on images of the atomic ensemble for the optimised MOT. Colour bars show OD scales, σ values show standard deviation of Gaussian fits to ensemble.

power ratios in the four MOT beams (approximately 1 (z): 0.8 (x): 0.4 ($y_1 + y_2$)), and therefore the relative confinement in these directions. A pair of rectangular MOT coils were again used, this time with more turns (50) and dimensions of 105×70 mm.

With this set-up we attempted to increase the length of the MOT during the compression phase by adding a shutter to the axial (y) beams. The idea was that we could optimise the power in the axial beams to obtain the highest atom number during the loading phase and then reduce this with the shutter just before turning off the MOT fields. This would reduce axial confinement and create a much longer MOT, and therefore achieve higher OD, in the y -direction. However, we were unable to see any improvement using this method and the shutter was removed from the experiment.

Combining the GEM and axial MOT fields also proved difficult. Initially two PBSs were used to combine and separate the fields. However, it was found that this did not allow for fine tuning of the GEM field polarisations and high efficiencies could not be achieved. Instead, we placed a small angle (1.5°) between the fields. This required moving all mirrors, lenses etc. for the GEM and MOT fields in the axial direction far from the MOT to reduce the angle as much as possible.

The MOT beam diameter in the x - and z -directions was 3.5 cm. This was not possible with the axial fields due to the minimum size of the glass cell down the y -axis, limiting the size to 2.5 cm. Further detail on this MOT set-up is shown in Fig. 12.16.

The methods used to optimise the first MOT were also used here. We employed a dark spot on the repump of approximately 3.5 mm. We also ramped up the magnetic fields from 20 to 40 G cm^{-1} during a 20 ms compression phase. During the same phase, the frequency of the trapping field was ramped down from -30 MHz to -70 MHz detuned and the repump from resonance to -8 MHz detuned. Atoms were pumped into the $F = 1$ level by turning off the repump $50 \mu\text{s}$ before the trapping field. If atoms were being imaged, then a $200 \mu\text{s}$ repump pulse was sent in immediately before the first imaging pulse to move atoms back to $F = 2$. An angle was again used for imaging down the line (I2), and an angle was now also required for side-on imaging (I1).

With these parameters we obtained the MOT shown in Fig. 12.6. These images were taken with an imaging beam detuning of -40 MHz. As can be seen, this MOT was significantly longer than the previous one, with an axial Gaussian standard deviation of




Figure 12.7: GEM Coil Switching. (a) Circuit diagram for switching coil: L - inductor (coil), 10-100 μH ; R - resistor, 1 Ω ; C - capacitor, 220 nF; M - MOSFET, model number IRF620. (b) Current in one coil, measured during magnetic field switching by placing a floating scope across the resistor, for (i) switch-off (blue) and (ii) switch-on (red).

$\sigma_y = 2.9$ mm. The larger MOT beams and powers also led to a larger total atom number of approximately $5 \cdot 10^9$ atoms. Combining these, we achieved a peak OD of approximately 250 on the imaging transition. The temperature of this MOT was higher than the previous one, with 150 μK in the y -direction, 460 μK in the z -direction, and 140 μK in the x -direction. The reason for the different temperatures in the x - and z -directions was most likely due to the magnetic field asymmetry in these directions.

It can be seen from Fig. 12.6(a) that the MOT was on an angle with respect to the horizontal. This issue could not be resolved by altering the axial optical fields or attempting to angle the MOT coils.

12.4.2 GEM Set-Up and Results with MOT 2

The lenses for the probe and coupling fields had to be changed from the previous set-up due to the much longer distance from the lens to the MOT. Using a much longer focal length lens for L2, the probe waist was 120 μm at the centre of the atomic ensemble and the coupling field was again collimated with a waist of approximately 500 μm . Due to the large distance from the cloud to the point where the GEM and MOT fields were separated, the coupling field was smaller than the probe before and after the MOT. To spatially filter the two in this case we used L3 and a pin-hole created by combining two slits of widths 200 and 100 μm . This allowed for 95% attenuation of the coupling field with 97% probe field efficiency. Again, with heterodyne detection, the remaining coupling field registered as a small dc offset on the signal.

With the new MOT orientation it was found that the method described in Sec. 12.3.2 for creating and switching magnetic fields no longer worked. This was because, with a larger separation needed between coils (70 mm) and larger radius (75 mm), the currents required to create the magnetic fields lead to inductive coupling between each pair of Helmholtz/anti-Helmholtz coils. Without this method we were left with the non-ideal




Figure 12.8: MOT 2 GEM Results. Log-log plot of the memory efficiency as a function of storage time for the MOT shown in Fig. 12.6. Points correspond to data averaged over approximately 100 traces, error bars show standard deviations, and the line shows a Gaussian fit to the data.

situation of using a pair of Helmholtz coils with large currents (2-4 A) and needing to raise a current in one coil and reduce it in the other. Apart from requiring that the gradient be switched quickly compared to the storage time, we also needed symmetric switching for the two coils and minimal oscillation after switching. Eventually, after much optimisation work, we created the switching circuit shown in Fig. 12.7(a), using a home-built actively-stabilised MOSFET-based switch system. With this set-up we achieved switching in $1.5 \mu\text{s}$ to within 1% of the desired current. This is shown in Fig. 12.7(b). More details on the switching circuit can be found in Ref. [426].

Though the peak OD measured down the line with the imaging system was much higher than in the previous set-up, the maximum OD measured for the probe was only 80. This was most probably due to the added difficulty in aligning the probe to the MOT with this set-up, as there was very little room for optimisation. Also, the angle of the MOT to the horizontal, coupled with the long distance the probe beam much travel between mirrors, made it difficult to have the MOT elongation and probe beam parallel.

As such, it was no surprise to find that the maximum efficiency in this case had dropped to 55%. However, the coherence time increased, with a Gaussian decay constant of $355 \mu\text{s}$, with 5% efficiency for storage over $500 \mu\text{s}$. This is shown in Fig. 12.8. This increase in coherence time with a hotter MOT indicates that there are other factors at play apart from temperature. For instance, for axial diffusion in GEM we are concerned with the ratio of ensemble length to temperature. In this case, the standard deviation had increased by a factor of two while the temperature increased by a factor of three. If we defined the length of the ensemble as being $4-6\sigma$ (i.e. 95-99.7% of the atoms) then the ratio would increase and we would expect to see an increase by roughly this amount in coherence time if we were limited by diffusion in the axial direction. Reduced magnetic field decoherence effects, either from the new GEM or MOT coil configurations, could also have played a role in increasing the coherence time here.




Figure 12.9: Cold GEM Set-Up 3, with: TR - trapping and repump beams; P+C - probe and coupling beams; 2D - 2D MOT coils (green); Ax - axial MOT coils (blue).

12.5 Configuration 3 - MOT Fields at 45° to GEM Fields, Again

From the previous set-up it was apparent that increasing the MOT length, and making it as thin and long as possible, greatly improved the MOT OD. It was also apparent that easy optical access for the GEM fields was important to achieve large probe absorption and therefore high efficiencies for the memory. This led to the development of our third, and final, MOT set-up.

12.5.1 MOT 3 Set-Up, Optimisation and Characterisation

This set-up was very similar to the original MOT set-up from Fig. 12.1. The main differences from the first set-up were those mentioned in Sec. 12.4.1: namely the use of the larger glass cell; increased power in the trapping fields; and a larger beam size of 3.5 cm for all three MOT beams. The size of the dark spot was increased to 7.5 mm. Further detail is shown in Fig. 12.17.

The other main difference was the MOT coil configuration. This is shown in Fig. 12.9. As can be seen, two different sets of coils were used to create the optimal geometry for the MOT. There was a set of 2D MOT coils, with dimensions of 70×150 mm and 25 turns, to create a spherically symmetric magnetic field around the z -axis, and a pair of axial anti-Helmholtz coils, with dimensions of 70×100 mm and 20 turns, to create 3D confinement. The shape of this MOT was determined by the currents in the 2D and axial coils, as well as the relative intensities of the trapping fields. In this case the ratio of trapping field power was 1 (y - z): 1 (z - y): 0.65 (x), with magnetic field gradients of 16 G cm^{-1} created by the 2D MOT coils and 2 G cm^{-1} for the axial coils in the loading phase.

We could collect over 10^{10} atoms with this configuration. However, we found that the static MOT parameters that optimised atom number did not optimise density in the

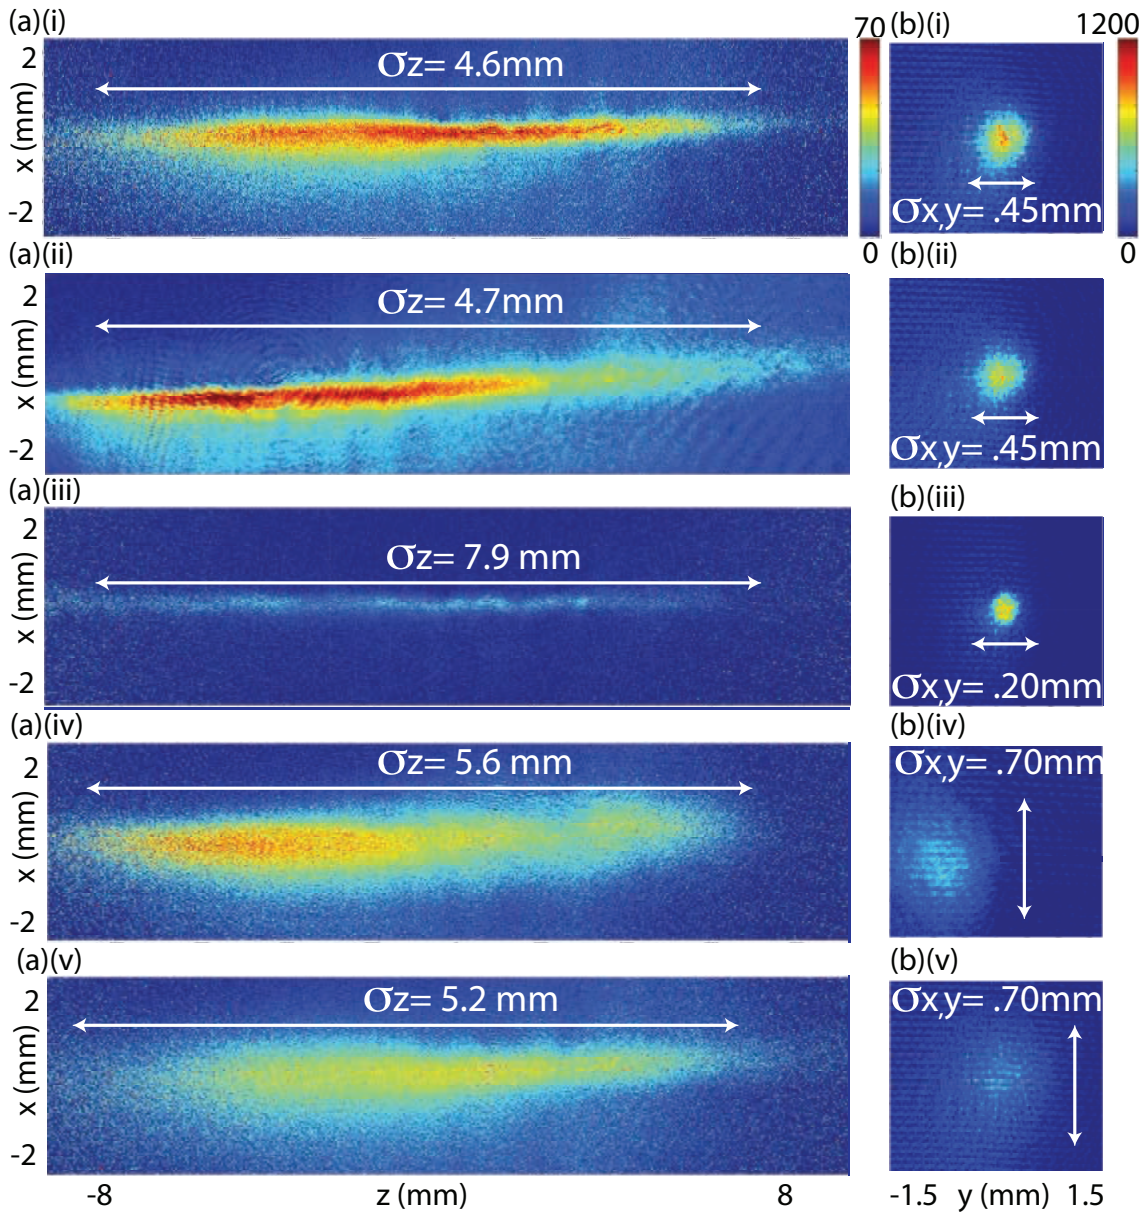


Figure 12.10: MOT 3 Characterisation Images. (a) Side-on and (b) front-on images of the atomic ensemble for various MOT parameters taken at $500 \mu\text{s}$ after the fields were turned off: (i) all optimisation protocols described in text; (ii) no dark spot; (iii) no axial gradient; (iv) no magnetic field compression; and (v) same as (i) but 3 ms after fields turned off. Colour bars on (a) traces show OD scales, σ values show standard deviation of Gaussian fits to ensemble.




Figure 12.11: Peak OD Characterisation. Averaged scaled (left-hand scale) and unscaled (right-hand scale) OD cross-sections, taken across 10 y -plane slices for seven traces, as a function of x position for an imaging beam detuning of (a) -60.8 MHz from resonance (OD scale factor = 410) and (b) $+59.2$ MHz (OD scale factor = 390). Blue points represent data, red lines represent Gaussian fits to the data.

compression phase, and we typically worked with a sample of $4 \cdot 10^9$ atoms.

In the 20 ms compression phase we ramped the optical frequencies identical to the previous configuration. In addition, the 2D magnetic field gradient in the x - and y -directions was ramped up to 40 G cm^{-1} . We did not ramp the axial field.

Images of the MOT for various configurations are shown in Fig. 12.10. For these imaging runs, we used a 4.98 s load time followed by 20 ms of ramping fields (described above). All fields were then turned off and an image of the MOT was taken $500 \mu\text{s}$ later (unless otherwise stated). As with the second configuration, a $200 \mu\text{s}$ repump pulse was used before imaging. For side-on imaging, the imaging beam was detuned by -20 MHz and by -60 MHz for front-on imaging to avoid complete absorption and, therefore, saturation of the measured OD, as well as diffraction effects. Unlike the first or second configurations, the front-on imaging beam was sent directly down the z -axis by placing mirrors before and after the MOT.

Figure 12.10 shows images of the atomic ensemble for various MOT parameters. The images in Fig. 12.10(i) were taken for a cloud with all the optimisation protocols described in this section and Sec. 12.4.1. The peak, front-on OD was measured to be 1000 on the $F = 2 \rightarrow F' = 3$ transition. This was achieved by detuning the imaging beam approximately 60 MHz for both positive and negative detunings to ensure no diffraction effects were present and using Eq. 11.2 to scale the OD. Two averaged OD cross-sections, taken across ten y -plane slices, for these detunings are shown in Fig. 12.11. The scaled peak ODs were 900 and 1100 for minus and plus detunings respectively. Averaging the maximum peak height of the individual traces gave over 1000 for both plus and minus detuning. The drop in the average cross-sections can therefore be attributed to slight movement of the MOT between images, with a narrower peak seen for the negative detuning. The temperature of this ensemble, determined by measuring the width of the expanding cloud 5, 10 and 15 ms after the fields were turned off, was approximately $200 \mu\text{K}$ in all directions.

Without the spatial dark spot, Fig. 12.10(a)(ii) looks very similar to Fig. 12.10(a)(i). However, the front-on image in Fig. 12.10(b)(ii) shows that the maximum OD dropped to 800 with $5 \cdot 10^9$ atoms. Figure 12.10(iii) shows the 2D MOT created without turning on




Figure 12.12: Probe Diffraction. CCD traces of probe absorption through the MOT for detunings of (a) -40 MHz and (b) $+40$ MHz. The colour bar shows log absorption scale.

the axial coils. This ensemble was much longer in the z -direction than the initial cloud (so long, in fact, that our imaging beam was not large enough to capture it all), but also contained only a quarter of the atoms. This meant that the peak OD dropped to about the same as that without the dark spot. The temperature in the x - and y -directions for the 2D MOT was approximately $100 \mu\text{K}$ but, as there was no trapping in the z -direction, the temperature here was much larger.

Figure 12.10(iv) shows the ensemble without ramping up the 2D MOT magnetic field at the end of the run. This led to a less-compressed cloud in all three dimensions, though mainly in the x - y plane, and therefore a much lower OD of approximately 500. Though there was less magnetic field compression, the temperature of this cloud was the same as the original cloud in the x - y plane, with $T = 200 \mu\text{K}$. However, in the z -direction the cloud barely expanded over 15 ms.

Finally, Fig. 12.10(v) shows the cloud from Fig. 12.10(i) expanded for 3 ms as opposed to $500 \mu\text{s}$. The temperature of this cloud should be the same as the initial cloud and, due to the expansion, the peak OD had fallen to 400, though the atom number was still around $4 \cdot 10^9$.

12.5.2 GEM Set-Up and Results with MOT 3

The experimental set-up for the third GEM configuration was similar to that shown in Fig. 12.1(c). The probe beam was more strongly focussed into the MOT to access the highest optical depths available in the atom cloud, with a waist of $50 \mu\text{m}$ at the centre of the MOT. The coupling field was again collimated, this time to a waist of 1.25 mm. More detail is provided in Fig. 12.17.

Using a CCD placed after a mirror at position Ps (see Fig. 12.1(c)) we investigated the probe beam transmission through the MOT as a function of detuning from the $F' = 2$ excited state. Two of these images are shown in Fig. 12.12. We observed that the high density of the atom cloud caused lensing of the probe beam, leading to changes in the image as a function of frequency. These CCD traces, combined with additional optical depth data from measurements made using an avalanche photo-diode at Ps and weak probe pulses, indicated a resonant probe optical depth of around 300 on the D1 $F = 1 \rightarrow F' = 2$ transition. Using the LO and heterodyne detection we performed the same experiment and measured an OD of 550. The difference between the two measurements was due to the change in mode shape seen in Fig. 12.12, which would translate to a lower visibility

on the heterodyne signal.

After the MOT we filtered the coupling field from the probe using lens L3 and the $100 \times 200 \mu\text{m}$ pin-hole (rather than the iris shown in Fig. 12.1(c)). In this case, the pin-hole removed over 99% of the coupling field and had a probe transmission efficiency of 90%. The heterodyne detection then saw no trace of the coupling field.

For the highest efficiency memory we used the MOT shown in Fig. 12.10(a). We again used a Gaussian pulse with a FWHM of $10 \mu\text{s}$, and the same switching set-up described in Sec. 12.4.2. For this pulse length, a one-photon detuning of -250 MHz , coupling field power of $350 \mu\text{W}$ (corresponding to a Rabi frequency of approximately 2 MHz), and peak probe power of 5 nW , we were able to demonstrate storage with $80 \pm 2\%$ total efficiency. This was measured by squaring the modulated pulse, finding the pulse shape, and integrating and averaging over 10 input and 17 output pulses, the error coming from the standard deviation of these traces. This is shown in Fig. 12.13(a), with the data being digitally demodulated in phase, averaged and then squared to produce the intensity plot. From the modulated heterodyne data we could also determine that the echo had the same frequency as the input pulse to within 100 kHz . As heterodyne detection is mode-sensitive, care was taken to optimise the visibility for the input pulse, so that any change in the mode during storage would lead to a reduction in the measured efficiency. During the experiment we observed that, when aligning the heterodyne position or polarisation, the input and echo increased or decreased together when the memory was working best.

To compare our memory to theoretical expectations, we needed to know about the decay of the information stored in the memory and the optical depth. To investigate the decay we delayed the gradient switching time to store the pulse for longer periods inside the memory, while turning off the coupling field during the storage window. This is shown in Fig. 12.13(b). For the MOT used to obtain the 80% efficiency above we found an exponential decay with a time constant of $117 \mu\text{s}$. The optical depth can be found directly from the storage data, where we observed 2% leakage of the probe field during the write phase. Assuming symmetric read and write operations, the total efficiency after $20 \mu\text{s}$ storage can be estimated as $0.98 \times 0.98 \times e^{-20/117} = 81\%$. We know from measurements of the broadened Raman lines (Fig. 12.13(c) - achieved by scanning the probe frequency around resonance with the coupling field on $500 \mu\text{s}$ after turning the MOT off²) that the write and read stages were not perfectly symmetric. However, the observed efficiency is still compatible with these measured bandwidths (between $50\text{-}150 \text{ kHz}$). We can also crosscheck this result using Eq.s 3.42 and 3.54 (normalising the bandwidth by γ). Taking the resonant optical depth of 150 (as only half the atoms reside in the m_F state used for the memory), the measured bandwidths of the read and write Raman lines and the decay time, we again estimate 81% total efficiency for our memory.

The decay found above was on a much shorter timescale than either of the other two MOT configurations, as well as having a different form. To further investigate the source of the decoherence we also measured the decay of the memory for two other MOT configurations. Firstly, we added a longer waiting period after switching off the MOT: 3 ms (Fig. 12.10(v)) instead of $500 \mu\text{s}$. This avoided any residual magnetic fields caused by eddy currents and resulted in an exponential time constant of $195 \mu\text{s}$, though with a lower initial efficiency due to the lower initial OD. As the temperature of clouds (i) and

²The attentive reader may notice that the absorption for the two Raman lines shown here will not give 80% efficiency when the exponential decay is taken into account. This is because they were taken on a different day to the 80% results, with sub-optimal alignment.

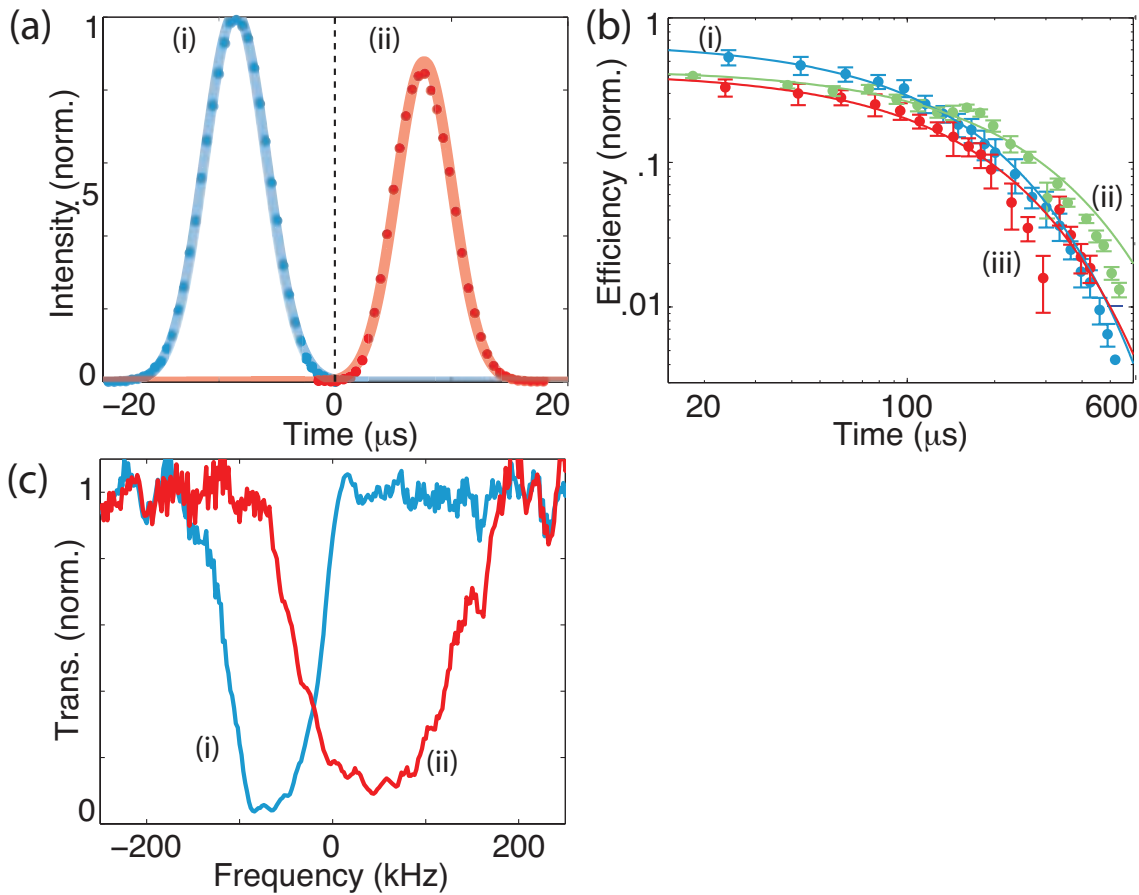


Figure 12.13: MOT 3 GEM Results. (a) High efficiency demodulated heterodyne data: (i) input pulse (blue), and (ii) echo (red) at 80% total efficiency. Points correspond to digitally demodulated data, averaged over 10 (17) traces for input (echo) and squared. Lines correspond to Gaussian fits. Dashed line indicates magnetic field switching at $t = 0$. (b) Log-log plot of memory efficiency as a function of storage time for: (i) MOT shown in Fig. 12.10(i) (blue); (ii) MOT shown in Fig. 12.10(v) (green); and (iii) MOT shown in Fig. 12.10(iii) (red). Points correspond to data averaged over approximately 10 traces, error bars show standard deviation, and lines show exponential fits to data. (c) Broadened Raman lines for the MOT shown in Fig. 12.10(a): (i) input (blue); and (ii) output (red) gradients using squared and averaged heterodyne data over approximately 20 traces, normalised to traces without a MOT.

(v) were the same, this indicates that temperature was not limiting the coherence time.

We also investigated the decay using a cloud without the axial coils (Fig. 12.10(c)). The initial efficiency was again lower than the first MOT, due to the lower OD and size of the cloud, and the time constant was 133 μs . This is interesting because, with no axial trapping, we would expect this cloud to decay much faster than the other two if the coherence time was limited by axial diffusion. The investigation of decoherence in this system is the subject of ongoing work, but could potentially be related to the tighter focus and corresponding smaller Rayleigh range of the probe for this set-up, or residual magnetic fields in any of the six MOT coils. Also, with the very high densities in the cloud, potentially interatomic interactions could be played a role in decoherence.

12.6 Discussion

While 80% total efficiency is, to the best of our knowledge, the highest efficiency so far reported for cold atomic ensembles, we have numerous paths for further improvement on this result. As with all atomic memory experiments, optical depth is a necessary condition for high efficiency. In our case, even with the large optical depth we have achieved, it is still a primary limit to our experiment, along with the decay rate. Though we only have 2% transmission for our 10 μs FWHM pulses, the decay rate will reduce the efficiency to at least 85%. Shorter pulses require larger bandwidths, which in turn require larger optical depths for similar efficiencies. It can be seen from the Raman features shown in Fig. 12.13(c) how an increase from approximately 50 kHz to 100 kHz bandwidth can quite severely reduce the absorption.

One limiting factor on OD is that, when pumped into the $F = 1$ state, approximately half the atoms end up in the $m_F = 1$ state used for the memory and half in the other two m_F states. If we were to implement the optical pumping scheme from Ref. [403] on the D1, rather than D2, transition we anticipate an increase in the OD by at least a factor of two. This would require the installation of two new lasers.

Another way to improve the density along the cloud would be to implement a dark line rather than a dark spot on the repump [401], as discussed in Sec. 9.4.2. By ensuring these lined up at the centre of the MOT, a high density cylinder rather than a sphere could be produced. Imaging the dark line onto the MOT using a lens would make it truly dark and improve this technique further.

Rather than increasing the density, we could also continue to increase the length of the ensemble by using a larger cell with larger (and therefore necessarily more powerful) trapping and repump fields. In addition, we have not yet tested our cold atom GEM memory on a sample of ^{85}Rb , which could potentially lead to a further factor of three increase in OD simply due to the larger relative abundance of the isotope if we are not limited by atomic density. A more drastic solution would be to move to a two-MOT configuration, as per Ref. [396], where we would expect up to ten times the number of atoms.

While we have demonstrated up to a seven-fold increase in coherence times over the warm vapour experiment, an important value for quantum repeater operation, the delay-bandwidth product (DBP) is also important. In this case, because we were using a pulse that was nearly a factor of five longer than in the warm vapour experiments, our DBP is approximately the same, if not slightly lower. To overcome this we need to either increase OD (as discussed above) to allow for larger bandwidths with the same high efficiencies, or

increase storage times further.

The storage time will be improved by removing sources of decoherence. One potential factor behind the decoherence in the final, ultra-high density MOT could be inhomogeneous background magnetic fields, partly due to eddy currents created by the switch off of the MOT coils. Improving this situation may require redesigning the switching electronics and the coil configuration. Atomic diffusion will clearly also play a role, so further reductions in temperature will be advantageous. One more extreme measure would be to transfer atoms into an optical lattice. The increase in coherence time afforded by such a measure must, however, be balanced against the likely decrease in atom number and thus storage efficiency.

Even without further improvement, our memory provides an excellent high-OD platform for numerous other proof-of-principle experiments. We could attempt the ac Stark gradient discussed in Ch. 10. Another interesting possibility is to investigate cross-phase modulation, discussed in Ch. 7, where the high OD, small interaction volume and long storage times are highly advantageous.

Another option would be to try to use the cloud to produce squeezing via polarisation self-rotation. This was attempted in a cold atomic ^{87}Rb ensemble by Horrom et al. [427]. They did not produce squeezing, however they had only 10^8 atoms and an OD of two. Finally, the high optical depth could be used for four-wave mixing to produce either entanglement [190] or single photons [150, 428].

12.7 Conclusion

In this chapter we have developed a cold atomic ensemble with $4 \cdot 10^9$ atoms and a peak optical depth of 1000 on the D2 $F = 2 \rightarrow F' = 3$ transition, specifically with quantum memory applications in mind. We used this ultra-high optical depth system to demonstrate the gradient echo memory scheme with a total efficiency of up to $80 \pm 2\%$ for pulses with a full-width-half-maximum of $10 \mu\text{s}$, still limited primarily by the optical depth. The decoherence of the final, high efficiency system was found to be exponential, with a time constant of 117-195 μs depending on the MOT parameters used, representing an improvement of a factor of two to four over the warm GEM experiment.

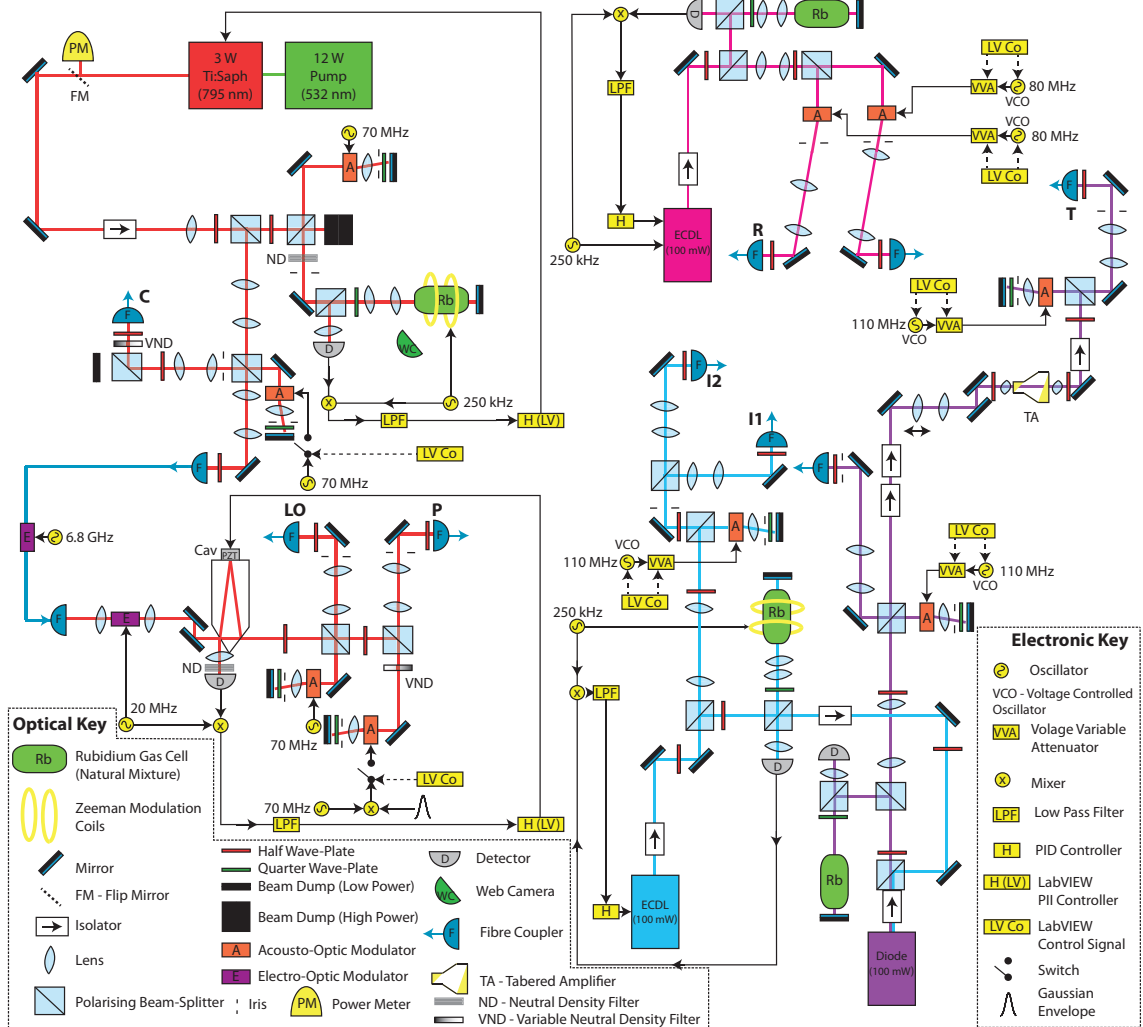


Figure 12.14: Full Cold Atom GEM Set-Up - Table 1. C - coupling field; P - probe field; LO - local oscillator field; T - trapping field; R - repump field; I1/2 - Imaging fields; Cav - optical mode-cleaner cavity; PZT - piezo-electric transducer. Frequencies of oscillators shown, LabVIEW control signals from triggering code, LabVIEW PII control provided by modified version of code described in Ch. 4.

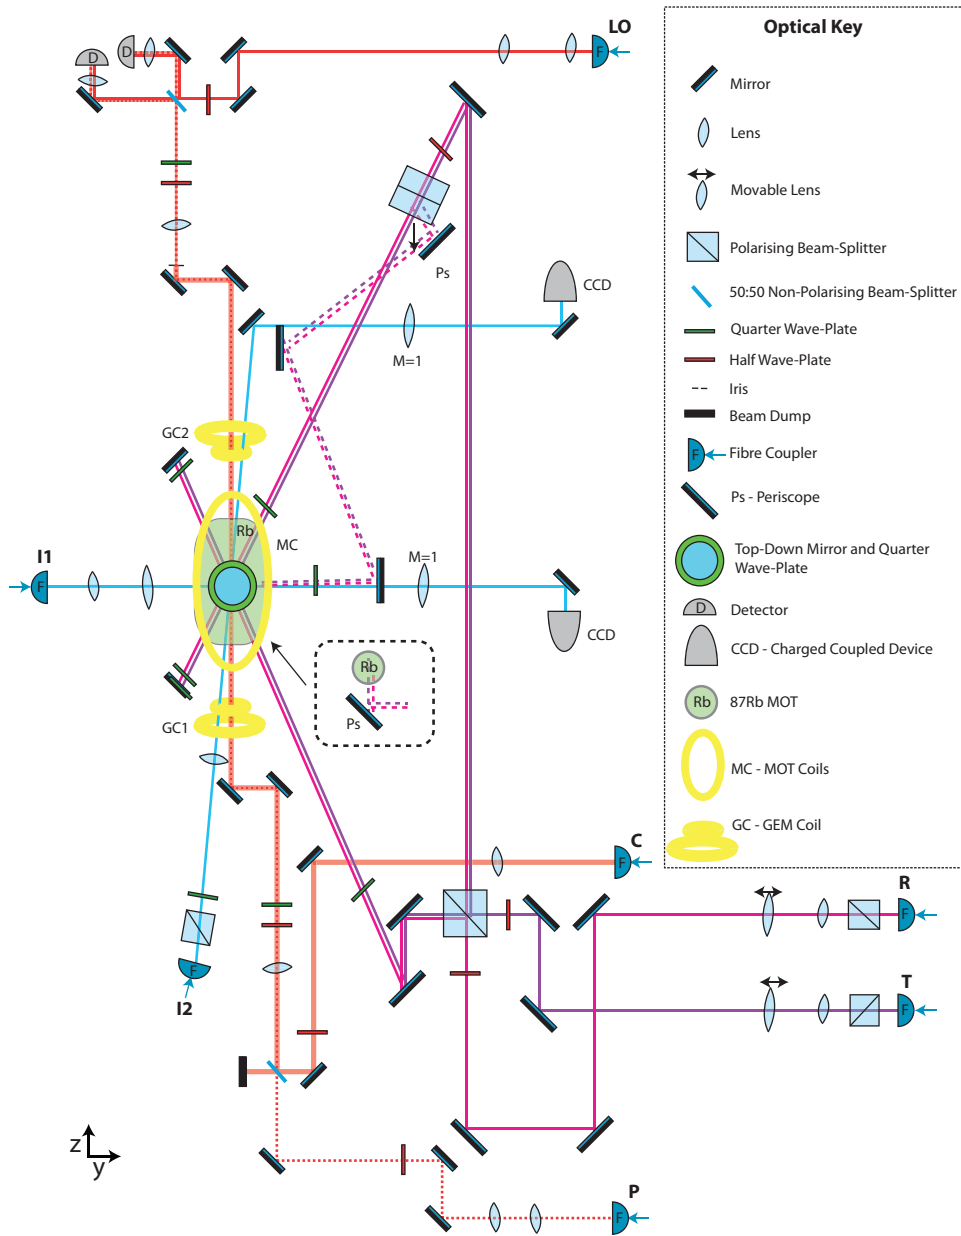


Figure 12.15: Full Cold Atom GEM Set-Up - Table 2.1. C - coupling field; P - probe field; LO - local oscillator field; T - trapping field; R - repump field; I1/2 - Imaging Fields. M values indicate magnification for imaging.

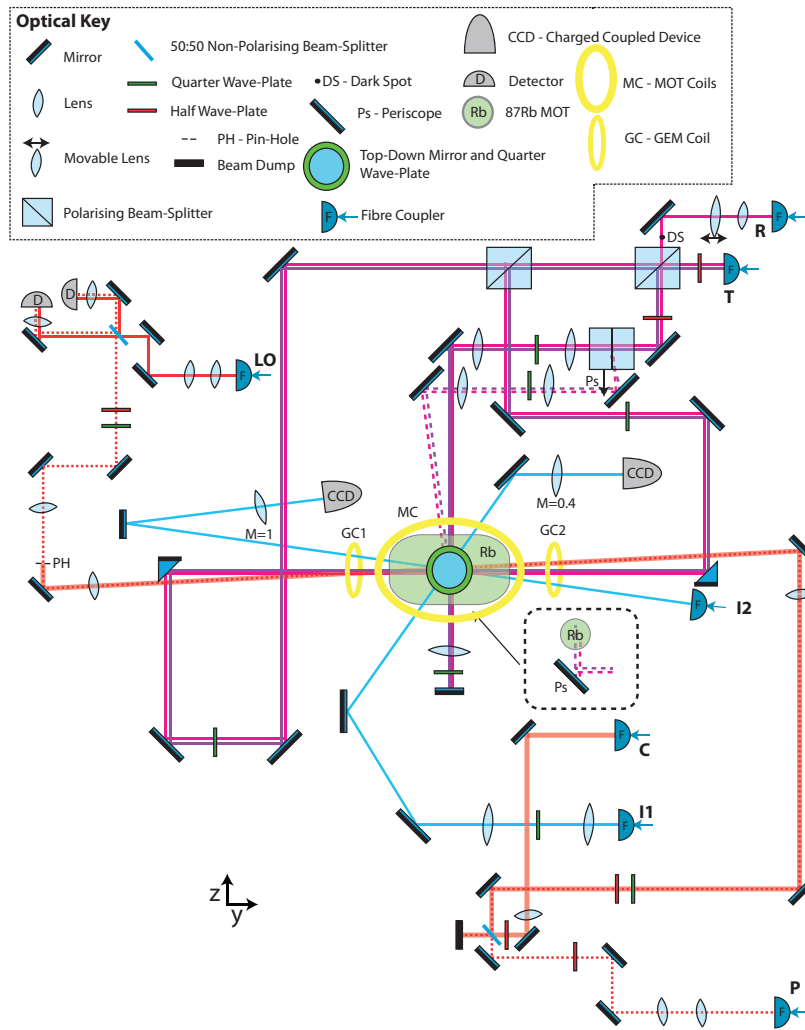


Figure 12.16: Full Cold Atom GEM Set-Up - Table 2.2. C - coupling field; P - probe field; LO - local oscillator field; T - trapping field; R - repump field; I1/2 - Imaging Fields. M values indicate magnification for imaging.

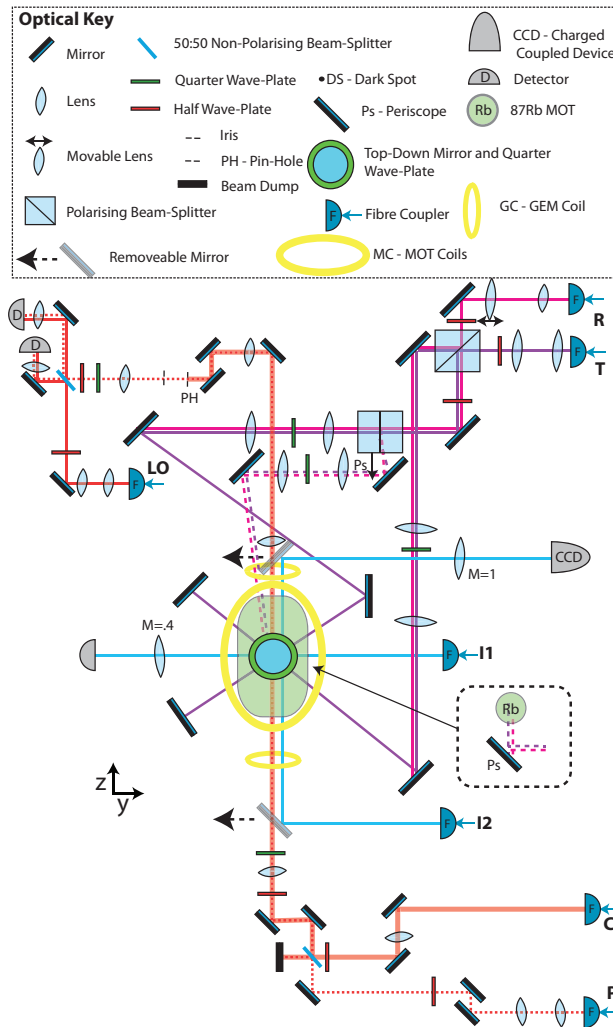


Figure 12.17: Full Cold Atom GEM Set-Up - Table 2.3. C - coupling field; P - probe field; LO - local oscillator field; T - trapping field; R - repump field; I1/2 - Imaging Fields. M values indicate magnification for imaging.

Conclusions and Future Outlook

Narrator: “At that moment... events occurred that are not, were not, and should never be considered an ending. For endings, as it is known, are where we begin.”
 - “Kerplunk” *Pushing Daisies*, ABC (America), 1 May 2009, TV Series.

In this thesis we have emphasised numerous times the importance of developing a quantum memory for quantum information processes. I hope by this stage the reader can appreciate that the lack of a highly-efficient, low-noise quantum memory with long storage times is a serious roadblock on the quantum information highway. I hope too that the reader can see the added functionality that being able to store multiple modes and manipulate stored information would bring to a quantum memory. Finally, I hope the reader will appreciate the delicate balancing act that is choosing a memory protocol and storage medium for a memory, with all their own advantages and disadvantages.

These points lie at the core of this thesis. Starting from a highly efficient and flexible memory protocol, in the form of the three-level gradient echo memory scheme, we have demonstrated that not only is the storage efficient but also noiseless, placing its operation above the no-cloning limit with a maximum unconditional fidelity of 98% for weak coherent states. The temporal multi-mode ability had already been demonstrated for GEM, with up to twenty pulse being stored and recalled [4]. We extended this to different spectral and spatial modes. The latter having the potential to increase bit rates in quantum repeaters by orders of magnitude and reduce the storage time restrictions [106, 105]. We have also made use of the frequency-encoding nature of GEM to spectrally manipulate stored information using a multi-element solenoid, with a view to performing qubit conditioning operations in large-scale quantum networks. We were able to precisely control the frequency or bandwidth of the echo, as well as interfere initially time-separated pulses of the same, or different, frequencies and perform a Fourier transform operation on the stored information. Manipulation of the phase of stored information was also demonstrated, using an external light field to produce a phase shift of 10^{-12} rad photon⁻¹ via cross-phase modulation, with a view to creating a single-photon controlled-not gate for linear-optic quantum computing applications.

The above-mentioned experiments were all performed in a warm rubidium-87 vapour cell. Vapour cells are an ideal test-bed for quantum memory experiments due to their ultra-high optical depths, ease of operation, and low cost. However, the movement of hot atoms was seen to be one of the main sources of decoherence for our memory, highlighted by the increased decay rate for higher-order spatial modes and images. The maximum coherence time achieved with this medium was 47 μ s. To attempt to push this value closer to the minimum millisecond-to-second regime required for quantum repeaters [100, 55] we decided to transport GEM to a cold atomic ensemble. This move required a

great deal of added complications to the set-up and it took multiple attempts and a few detours (such as experimentally characterising the transition between electromagnetically-induced transparency and Autler-Townes transmission windows) but the results speak for themselves: an ensemble of atoms with a peak optical depth of 1000 and a temperature of approximately $200 \mu\text{K}$; a memory efficiency of $80 \pm 2\%$ - the highest for a cold atom memory and comparable with the highest achieved with warm GEM of 87% [4]; and coherence times ranging from 117-355 μs , factors of two to seven times longer than in the warm vapour cell. The high densities and small volumes achievable with cold atomic ensembles also offer us an alternate gradient creation mechanism in the form of the ac Stark shift [367, 366]. Not only would this style of gradient improve switching times by potentially orders of magnitude over the current magnetic field gradients by using Pockels cells, but spatial-light modulators could also be used to further increase the precision control of the gradient and therefore improve spectral manipulation operations. A bandwidth of 150 kHz W^{-1} was calculated for the ac Stark gradient, meaning bandwidths on the order of megahertz should be possible with physically reasonable laser powers.

Another theme running throughout this thesis was the use of digital control to help operate complex quantum optics experiments. The capstone of this work was the development of a digital locking code. The code's ability to implement sequential automatic locking logic reduces the workload of the user, allowing them to focus on their ultimate goal as opposed to the cavities and phase-locks that are the nuts and bolts of most quantum optics experiments. The ability of such a system to collect and display all relevant information in one location was used, for instance, to develop a lock analyser tool. The flexibility of digital control was demonstrated in most chapters, most notably for the acquisition of data for characterising Schrödinger cat states, but has now been implemented on multiple experiments in Canberra, through to Brisbane [429] and Tokyo [430].

The pinnacle of the work presented here was the development of the cold atomic source for GEM, as it holds the promise for improving nearly all the other facets of GEM investigated in this thesis: an ac Stark gradient along with the cold atoms could improve the precision of the spectral manipulation experiments (though more work is required to efficiently shape the light gradient); the ultra-high atomic densities and low temperature could lead to a significant single-photon cross-phase modulation; the low atomic velocities should also improve the image storage abilities of the system. The 80% efficiency results achieved with the cold atom system are all the more exciting because there is still room for optimising: optical pumping alone should increase the OD by a factor of two; while moving to rubidium-85 or a two-component MOT could also improve efficiencies. The ultra-dense atomic source created also has potential uses in other quantum optics experiments, such as four-wave mixing to produce squeezing or single-photons. Similarly, the fast magnetic field switching circuit designed for the cold atom experiment could find uses on the warm vapour cell experiment to further improve switching behaviour. Coherence times could be improved by moving to a dipole trap, though not without much experimental work to ensure optical depths do not drop significantly, especially if a standing-wave dipole trap is to be used. Sources of decoherence apart from transverse diffusion, such as magnetic field decoherence, inter-atomic collisions and longitudinal diffusion, will also need to be investigated in more depth to push storage times over the millisecond mark.

For GEM, the final frontier is now the storage of a truly quantum state, such as single-photons or squeezing/entanglement. Due to the lower coupling field powers and noise in cold atomic systems, the cold GEM set-up lends itself to single-photon storage. The small Doppler-broadening also makes cold atom GEM suitable for memory-based

linear quantum computing protocols [431].

Alkali Structure

Rainer Wolfcastle: “Up and at them.”
Dialogue Coach: “No, ‘Up and atom’.”
 - “Radioactive Man” *The Simpsons*, Fox, 24 September 1995, TV Series.

In this appendix we briefly summarise the structure of the two atoms used in the thesis, namely rubidium-87 and caesium-133. Values presented below are from Ref.s [184] and [432]. But first a few notes on quantum-mechanical atoms.

A.1 Comments on Level Structures

The level structure for any atom depends on the number of protons and electrons Z , the nuclear and electron spin (\mathbf{I} and \mathbf{s} respectively), the orbital angular momentum \mathbf{l} , vector sums $\mathbf{J} = \mathbf{s} + \mathbf{l}$ and $\mathbf{F} = \mathbf{J} + \mathbf{I}$, and their interactions $\mathbf{J} \cdot \mathbf{I}$ and $\mathbf{l} \cdot \mathbf{s}$ [204]. Here we focus on the different F levels of the alkali atoms in question. Each F level will contain $2F + 1$ m_F hyperfine states. These will all be degenerate without the application of electric, optical or magnetic fields. When these fields are applied they are known by the names dc Stark, ac Stark and Zeeman shifts respectively. As we focus entirely on alkali atoms in this thesis, and as these atoms have no dc Stark shift [404] and an entire chapter is dedicated to the ac Stark shift (Ch. 10) we will quickly mention the form of the Zeeman shift here.

In the presence of a magnetic field with magnitude $|B|$ the m_F states of the F levels will change in energy by an amount [204]

$$\Delta E = g_F \mu_B m_F B, \quad (\text{A.1})$$

where g_F is the Landé g-factor and μ_B is the Bohr magneton. If we combine g_F and μ_B we can determine the splitting as function of MHz G^{-1} and it is these numbers that are presented below. The above equation assumes that $\mu_B |B|$ is smaller than the hyperfine state splitting (see below).

There are also only certain allowed transitions between m_F states, namely $\Delta F = 0, \pm 1$ and $\Delta m_F = 0, \pm 1$. The change in m_F state will be determined by the polarisation of the input light with π -polarised light leading to $\Delta m_F = 0$ and σ^\pm light leading to $\Delta m_F = \pm 1$.

A.2 Rubidium

There are two natural isotopes of rubidium; rubidium-85 (^{85}Rb , 72%) and rubidium-87 (^{87}Rb , 28%). In this thesis we will use ^{87}Rb for our memory. ^{87}Rb , as with all alkali atoms, has a well known structure with, in this case, two hyperfine ground states $S_{1/2}, F = 1, 2$




Figure A.1: Full Rubidium-87 Level Structure. (a) D1 line and (b) D2 line with: g - m_F -state Zeeman splitting ($g = .93$ MHz/G for all $P_{3/2}$ levels); d - dipole moment, ω - frequency of the transition; λ - wavelength of the transition; γ - excited-state decay rate; Δ_{hfs} - hyperfine splitting between ground-state F levels; Δ'_{hfs} - hyperfine splitting between excited-state F' levels.

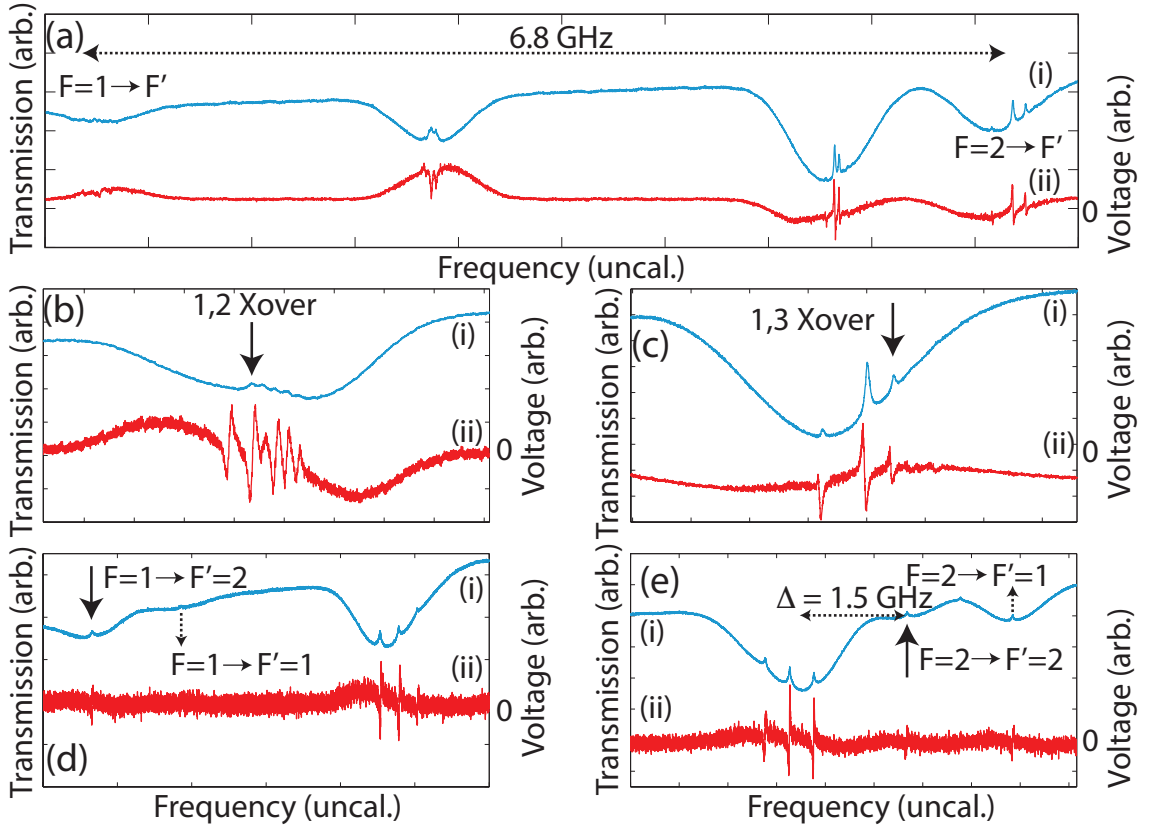


Figure A.2: ^{87}Rb Saturation Absorption Traces, with uncalibrated frequency scales for (a) full D2 transition; (b) close-up of D2 $F = 1 \rightarrow F'$ transitions (repump); (c) close-up of D2 $F = 2 \rightarrow F'$ transition (locking); (d) close-up of D1 $F = 1 \rightarrow F'$ transitions and (e) close-up of D1 $F = 2 \rightarrow F'$ transitions for (i) transmission (blue, left hand scale) and (ii) error signals (red, right hand scale). Non-labeled peaks are due to ^{85}Rb present in the vapour cells, solid arrows indicate locking points (Xover = crossover), dashed arrows indicate other peaks of interest (Δ is the one-photon detuning used for warm atom Λ -GEM experiments).

and both a D1 and D2 transition, which correspond to $P_{1/2}$ and $P_{3/2}$ levels respectively with wavelengths of 795 nm and 780 nm. This structure is shown in Fig. A.1, containing the hyperfine structure splittings Δ_{hfs} and $\Delta'_{hfs}(J')$ for the ground and excited states respectively, as well as the Zeeman splitting for all levels and on-resonance dipole moments. The fine structure splitting $\Delta_{hf} = 2\pi \cdot 7$ THz between the $P_{1/2}$ and $P_{3/2}$ transitions is not shown. Of particular interest to us is the Zeeman splitting of the ground states ± 0.7 MHz G^{-1} , as this will determine the size of the frequency gradient we apply across the memory to be 1.4 MHz G^{-1} . This comes about as we use the same m_F state in both F levels. It is also worth noting again that ^{87}Rb does not have a dc Stark shift [404], which is why magnetic fields were used to avoid incredibly large and unwieldy electric fields (on the order of 300 kV MHz^{-1}).

Saturation absorption spectroscopy is a useful tool for being able to view the hyperfine structure of atoms in warm vapour cells. It is also very useful for locking to atomic transitions (see Sec. 11.2.1). Traces of the relevant ^{87}Rb lines used for locking lasers are shown in Fig. A.2.

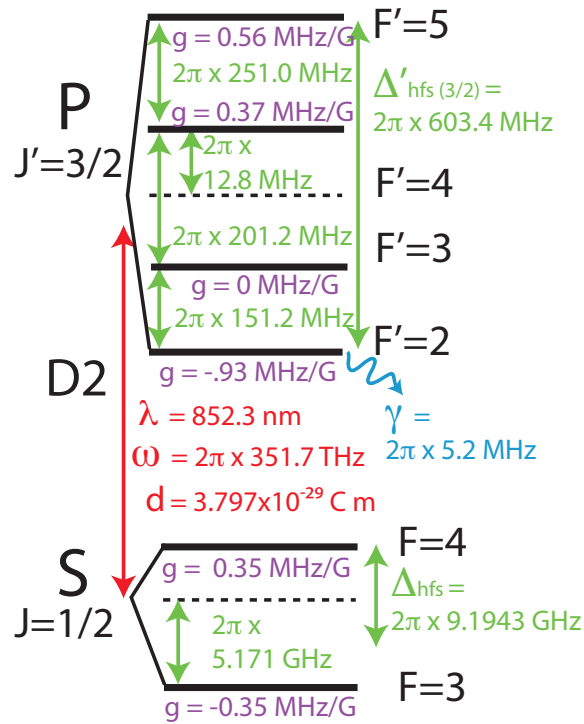


Figure A.3: Caesium-133 D2 Level Structure, with g - m_F -state Zeeman splitting; d - dipole moment, ω - frequency of the transition; λ - wavelength of the transition; γ - excited-state decay rate; Δ_{hfs} - hyperfine splitting between ground-state F levels; Δ'_{hfs} - hyperfine splitting between excited-state F' levels.

A.3 Caesium

The structure of caesium-133 is similar to that of ^{87}Rb . We only use the D2 transition in this thesis, and the structure of the relevant levels, detunings and m_F splittings are shown in Fig. A.3.

PDH Lock 5001 Instruction Manual - Rev. 0.1

Cyborg: "You can't mess with the program, Ralph!"
- *Wreck-It Ralph*, Walt Disney Pictures, October 2012, Movie.

This code is released under the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Licence (<http://creativecommons.org/licenses/by-nc-sa/3.0/>). If you do use the code and feel that it has contributed significantly to any work accepted for publication we ask that you acknowledge the following work:

- B. M. Sparkes, H. M. Chrzanowski, D. P. Parrain, B. C. Buchler, P. K. Lam, and T. Symul, *A scalable, self-analyzing digital locking system for use on quantum optics experiments*, Review of Scientific Instruments **82**, 075113 (2011).

We are interested in cataloguing modifications made to the code. If you do modify the code in a substantial way we would appreciate it if you would notify us and send us a copy of the code so that we can place this on our website for others to download.

If you have any comments, suggestions or queries please send an e-mail to ben.sparkes@anu.edu.au.

Copyright © 2011 B. M. Sparkes, All rights reserved.



B.1 Introduction

The goal of this work is to provide an easily modifiable code to lock multiple cavities and interferometers for quantum optics experiments. This code uses the Pound-Drever-Hall locking technique [246, 247] to allow for automatic re-locking. The code also includes a White Noise Generator that can be used to investigate the frequency response of the locks, as well as sequential locking logic. For more detail on the background and motivation for this work, please see Ref. [433] attached.

B.2 Set-up

This code requires National Instruments LabVIEW 2010® (32 bit) to run. Other requirements and basic set-up are detailed below.

B.2.1 Hardware

Below is a list of equipment that has been used while developing the *PDH_Lock_5001* code. The bracketed number indicate the numbers required for 8 locks.

- PXI chassis - NI PXI-1042Q with PXI-8106 embedded controller¹;
- FPGA cards - NI PXI-7852R (x2);
- Frequency generator - NI PXI-5404 (x1);
- AIP/AOP break out box (BoB) - made in house, 8 analog inputs and 8 analog outputs using Connector 0 line from FPGA (x2), also FPGA 1 needs to have 8 digital output lines (DIO0-7) connected for clock programming protocols (see below);
- High speed ADC - AD9460BSVZ-80 (x8);
- Clock generator - AD9959/BZ (x2).

Figure B.1(a) shows the basic set-up for a single lock. In all cases the Freq. Gen. signal (at 80 MHz) must be split (in our case using a MiniCircuits splitter, model number: ZSC-2-1), with one output being sent to the clock generators and one to the HS ADCs, both of which are split again as required². The modulated signal from the detectors (normally on an ac output) are sent one to each HS ADC. As each FPGA card has only two digital connectors (1 and 2), to send four digital signals from the ADCs to one FPGA requires combining the signals from two ADCs into one connector. In the code it is assumed that the first ADC of each pair is wired to lines DIO0-15 of the relevant connector and the second is wired to lines DIO16-31.

The low speed AIP/AOPs do not require clocking and will run between 1 M sample s^{-1} and 750 k samples s^{-1} . Normally 4 AIPs and 5 AOPs are used per BoB (AIP0-3, AOP0-4): the four inputs for the DC signal; and the first four outputs (AOP0-3) for the Scan/Lock signal which are, in turn, sent to the High Voltage amplifier. AOP4 carries a scan trigger for external hardware (see Section B.7.2 for more details). Also, at least

¹Only 1 chassis and controller required for 16 locks.

²NB: To avoid clock phase issues between HS ADCs ensure that the same cable length and splitters are used for this secondary - as well as tertiary and quaternary - splitting. See Section B.8.

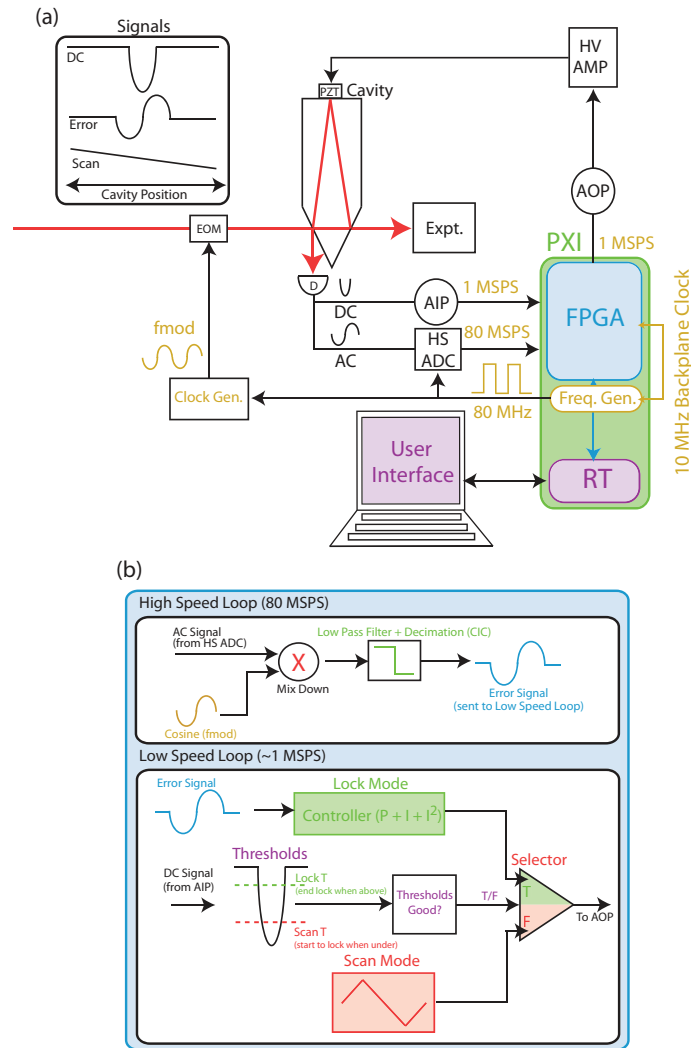


Figure B.1: Set-Up. (a) Depiction of the physical system and hardware required for digital locking using this code, in this case for a mode-cleaner cavity. PXI - PXI chassis, FPGA - field programmable gate array card; Freq. Gen. - frequency generator; RT - Real Time controller; AIP - analog input (sampling frequency $< 1 \text{ M sample s}^{-1}$); AOP - analog output (sampling frequency $< 1 \text{ M sample s}^{-1}$); HS ADC - High Speed analog to digital converter ($80 \text{ M samples s}^{-1}$); D - detector; Clock Gen. - digital clock generation board; f_{mod} - modulation frequency for lock; EOM - electro-optic modulator; PZT - piezo-electric transducer; HV Amp - high voltage amplifier; and Expt. - rest of experiment. Inset shows the different signals which are used for locking: DC - dc transmission/reflection signal; Error - error signal; Scan - scan function (provided by FPGA code). (b) Depiction of LabVIEW® 2010 code used to program FPGA cards. In the high speed loop (running at $80 \text{ M samples s}^{-1}$) the input signal from the HS ADC is mixed down with a cosine signal at f_{mod} (generated by a look-up table in LabVIEW®). The signal then passes through a low pass filter (LPF) and is decimated to produce the error signal. This error signal is, in turn, passed to the low speed loop (running at $\approx 750 \text{ k samples s}^{-1}$) where, depending on the value of the DC signal and the threshold limits imposed by the user, the system will move from Scan Mode, where the cavity is scanned using a sawtooth function, to Lock Mode, where a Proportional-Integral-Double Integral (PII) controller is used to keep the system on resonance with the laser. The code will move from Scan to Lock mode when the DC signal moves below the user defined Scan Threshold and will then remain in Lock mode until the DC signal rises above the Lock Threshold. If this occurs, the cavity will Scan again until the system returns to resonance, etc.

the first (master) FPGA BoB needs to have access to 8 digital lines. These are used to program two AD9959 DDS (i.e. clock generation) boards. Six digital lines (DIO0-2,4,6-7) are sent to all boards, with DIO5,3 sent to DDS boards one and two respectively (for more information on AD9959 programming please download the relevant documentation from the Analog Devices website [434]). The outputs from the Clock Gen. box are sent to the related modulators.

B.2.2 Software

As mentioned above, this code was developed using LabVIEW 2010® (32 bit) and therefore this version, or better, is required to run it. An overview of the FPGA code is shown in Figure B.1(b), and consists of two loops: a high speed loop running at 80 M samples s^{-1} ; and a low speed loop running at approximately 750 k samples s^{-1} (limited by the AIPs/AOPs). In the high speed loop the modulated signals from the HS ADCs are demodulated with a cosine function generated using a look-up table at the same frequency as the modulations (i.e. f_{mod}). These pass through a Cascaded-Integrator-Comb (CIC) filter [248] to produce error signals which are then sent to the low speed loop.

There are two main components to the low speed loop: a Lock Mode (consisting of a P, I and I^2 controller); and a Scan Mode (consisting of a sawtooth scan function). If the code is set to **AutoLock** the code will move from the Scan Mode to Lock Mode depending on the thresholds set for the DC signal: the code will start to lock when the DC signal dips below the Scan Threshold and will stay in Lock Mode until the signal rises above the Lock Threshold. If the lock has no DC component (i.e. a Michelson-style lock) the absolute value of the error signal can be used to set the thresholds instead (see Section B.4 for more information). The low speed loop includes a white noise generator (with the output also being sent to the same AOP as the Scan and Lock signals).

The *Master* and *Slave* FPGA codes differ in that the *Master* code contains a trigger for the Freq. Gen. (to try to minimise the issue of the ADC and FPGA running out of sync - though if this occurs then it can be rectified by varying the clock phase, see Section B.7.2) and also code to run two Clock Gen. boards. The *Slave* code does not contain this extra logic. For more than 8 locks the *Master* code should be compiled on the third FPGA (minus the Freq. Gen. triggering code) and the *Slave* code on the fourth.

The Real Time code acts as the interface between the FPGAs and the user, as well as to allow communication between FPGAs for sequential locking purposes (see Section B.7.3). It also controls the Freq. Gen. cards. Only one RT code is required for all FPGAs on one chassis. Included in the *PDH_Lock_5001* project are RT codes for 4 and 8 locks. These can be extended to 12 or 16 without much additional coding.

B.3 First Time

The following instructions are for when opening the project for the first time. Before running for the first time make sure all components and drivers for LabVIEW® 2010 are installed.

1. Open the *PDH_Lock_5001* folder, select the *PDH_Lock_5001.lvproj* project file.
2. In the project folder right click on the *RT PXI Target* and select *Properties*. In this menu enter the correct name and IP address of your PXI chassis (if unsure check using Measurement and Automation Explorer - MAX).

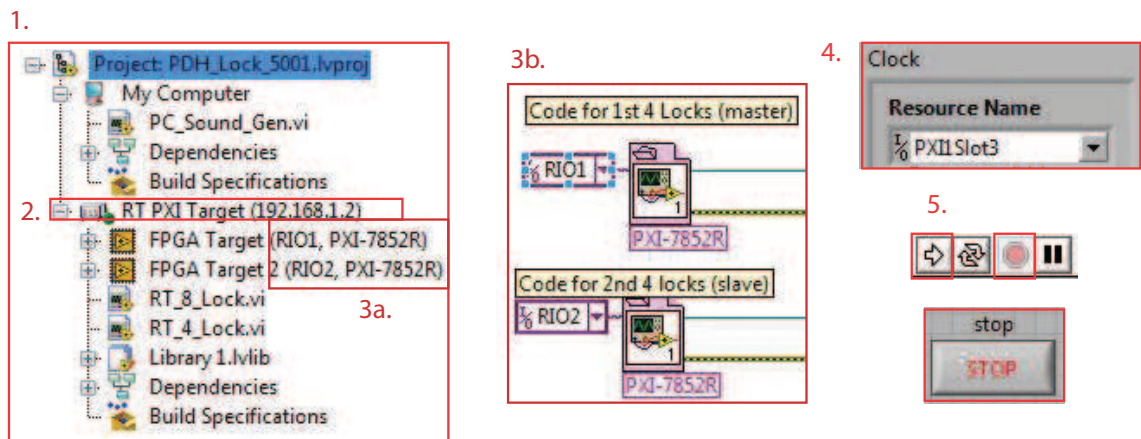


Figure B.2: First Time. The menus and controls which need to be altered before using the code for the first time. The numbers and boxes refer to the instructions in the text.

3. Make sure that the RIO# of the FPGA cards are correct (check using MAX). To change, right click on the relevant *FPGA Target* and select *Properties*. If you do change the RIO# here you will also have to alter this in the RT codes by opening the RT code, using Ctrl E to bring up the Block Diagram, scroll to the far left hand side of the code and change the Dialog box before the relevant FPGA.
4. Check the **Resource Name** of the clock (located in the **Clock and Timing and Misc** tab of the **System Inputs/Outputs** folder) against that in MAX. Change if different.
5. To run the code select the white right arrow icon in the top left hand corner of the Front Panel. To stop the code use either the red octagon next to the start button (this is for force quit) or, better, use the **Stop** button included in the code located above the **System Inputs/Outputs** folder.

B.4 Locking the System

Below are listed the suggested steps to achieve lock using the *PDH_Lock_5001* code (assuming the hardware has been set-up correctly). An overview of the various menus and more in depth description of the various controls is included in Section B.7.

1. Make sure that the Scope (Located in the **Graphs** menu) is set to the correct lock using **Scope Trigger Channel** and the Scope is turned on (located in the **Scope** tab of the **system Inputs/Outputs** menu), this will allow you to see your progress on the computer screen (see Figure B.3). If the scope is not displaying or triggering correctly try increasing/decreasing the **Acquisition Rate (kHz)** and **Scope Time Out (ms)** controls (not shown).
2. On the **System Inputs/Outputs** menu select the **Lock #** tab of interest and then choose the modulation frequency of the lock via the **Freq (MHz)** control in the **Signal** cluster. This automatically changes the frequency sent to the relevant EOM

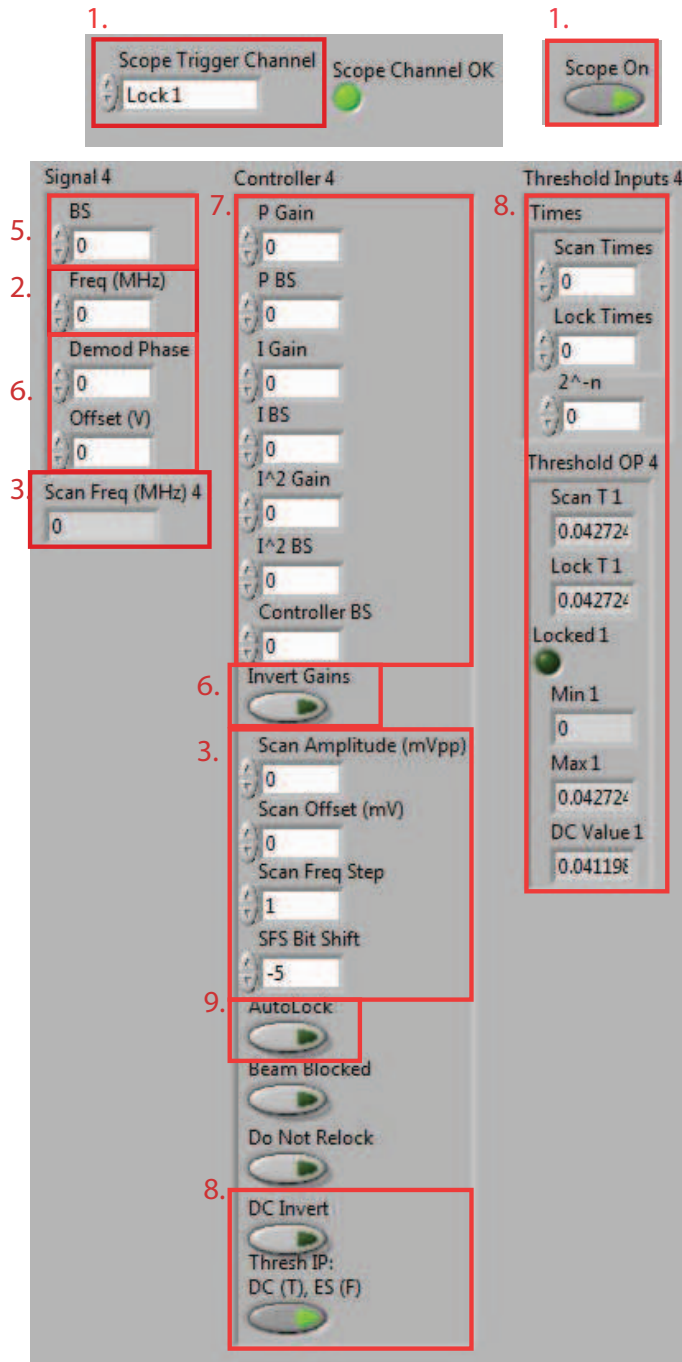


Figure B.3: Locking the System. The menus and controls used to lock the system. The numbers and boxes refer to the instructions in text.

by the DDS board as well as the demodulation frequency in the code³.

3. Set the **Scan Amplitude (mVpp)**, **Scan Offset (mV)** and **Scan Freq Step and SFS Bit Shift** to achieve a reasonable scan range and rate (see **Scan Freq (MHz)** indicator), all controls are located in the **Controller** cluster. The range of the AOPs is $-10\text{ V} \rightarrow 10\text{ V}$.
4. An error signal should now be visible of the scope. If not, check that the correct modulation is being sent to the EOM and that the right input signal is being received by LabVIEW®. To check if the input is working properly, try sending the modulation signal straight to the HS ADC input, which should produce a straight line on the Scope.
5. Once an error signal has been produced, use the **BS** control in the **Signal** cluster to increase/decrease the size of the error signal by factors of 2 (to avoid saturation).
6. Maximise the size of the error signal using the **Demod Phase** control. If the error signal is 180° out of phase (see Figure B.1(a) inset), use the **Invert Gains** control to avoid having to use negative gains. If there is a constant DC offset when the phase is maximised, use the **Offset (V)** control to remove it.
7. Set some gain for the **P Gain** and **I Gain** components of the controller, as well as the **Controller BS** (normally starting around -15)
8. Set the **Scan Times** and **Lock Times** values ($\text{Scan Times} \leq \text{Lock Times}$), as well as the **2ⁿ** control to set the scan and lock thresholds (see Section B.7.1). Make sure that, if there is a DC signal for the lock, it is displayed on the scope. If not, then toggle the **Thresh IP: DC (T), ES (F)** switch on. If it is a Michelson-style lock, with no DC signal, the absolute value of the error signal can be used instead by turning the **Thresh IP: DC(T), ES(F)** switch to false. Also check that the DC signal is the correct way up (see Figure B.1(a) inset), if not then use the **DC Invert** control. The maximum, minimum and current values of the DC input, as well as the threshold values, are shown in the **Threshold OP** cluster. Set the **I1 Max Value** and **I2 Max Value** thresholds in the **Clocks and Timing and Misc** tab (not pictured). This will cause the locks to reset if the I or I² controller values rise above these levels. The I and I² values can be viewed in the **Re-Ordered IP/OPs** folder in the **I Values** tab.
9. Set the thresholds using the **Threshold Inputs** cluster. The Scan and Lock thresholds are fractions of the size of the DC signal determined by **Scan Times** $\times 2^{-n}$ and **Lock Times** $\times 2^{-n}$. The higher the fraction, the closer to the bottom of the DC signal the threshold will be. The system will lock once the DC value goes below the Scan Threshold and remain locked until it rises above the Lock Threshold, therefore ensure that **Scan Times** \geq **Lock Times**
10. Turn on the **Autolock** control. The system will do one full scan to determine the maximum and minimum DC values and calculate the Scan and Lock thresholds (see **Threshold OP** cluster) before trying to lock (see previous step). If the system does not lock, make sure that the phase is correct (see Figure B.1(a) inset for correct

³To change the modulation depth of the signal, use the relevant **Amp IP Cluster** for the DDS board located in the **Clocks and Timing** tab, see Section B.7.2.

orientation for positive gains), if this is not the case try the **Invert Gains** control. If this does not work, then try decreasing (or increasing) the gains, decreasing the scan frequency, the scan amplitude, or decrease the scan and lock times. You may also want to check the dependencies of the lock to make sure that it is not **Beam Blocked** (see Section B.6). Once locked, to optimise the gains see Section B.5.2.

11. As long as **Autolock** is on then the system will either be locked or trying to find lock. If lock is dropped and the system cannot satisfy the set Scan and Lock thresholds (due to power drifting of the laser etc.) the maximum and minimum values will be retaken after a time **Lock TimeOut (s)** located in the **Clock and Timing and Misc** tab. If it is not desirable for the system to relock then select the **Do Not Relock** control. In this case, once the lock comes unlocked the code will move back to Scan Mode until it is turned off. Also, to draw attention to this, if the *PC_Sound_Gen.vi* code is run in parallel with the RT locking code, a sound will be made when a lock has come unlocked and **Do Not Relock** is active for it.
12. Finally, if you desire to stop the lock exactly where it is and freeze the maximum and minimum DC values (for instance if you will be blocking the beam before the lock), select **Beam Blocked**.

B.5 Lock Optimisation

In this section we will go through a brief explanation of how the code is designed to allow for lock optimisation.

B.5.1 Control Theory

Figure B.4(a) shows the basics of a closed loop feedback system, as well as the definitions we will use in this section⁴. The controller transfer function H , used for calculating GH which determines the system bandwidth and phase margin, is a simulated H using the following equation [244]

$$H_{sim}(z) = k_P + k_I \cdot \frac{z}{1+z^2} + k_{II} \cdot \left(\frac{z}{1+z^2} \right)^2, \quad (\text{B.1})$$

(where k are the gains, $z = \exp(iff/f_s)$, and f_s is the sampling frequency of the controller system). Another way of determining H is by calculating it through G and T , $H_{calc} = (G - T)/(GH)$, however this does not work once the controller gain drops below a certain level (close to unity gain).

A Network Analyser (shown in green) can be used to measure the magnitude ($20 \cdot \log_{10}|T|$) and phase plots of the system. In our system we use a white noise generator function contained in the LabVIEW® code and average over many cycles, whilst a store bought network analyser⁵ usually uses a swept sine wave. Figure B.4(b) shows a comparison of these methods with good agreement at higher frequencies (≥ 700 Hz). There is some statistical noise at low frequencies for the LabVIEW® white noise generator, but this can be improved by increasing the number of averages taken for measurements. Measuring

⁴Note that we define the system bandwidth to be the unity gain frequency, and the phase margin follows from the Nyquist stability criterion [245].

⁵Here, a Model MS4630B from Anritsu.

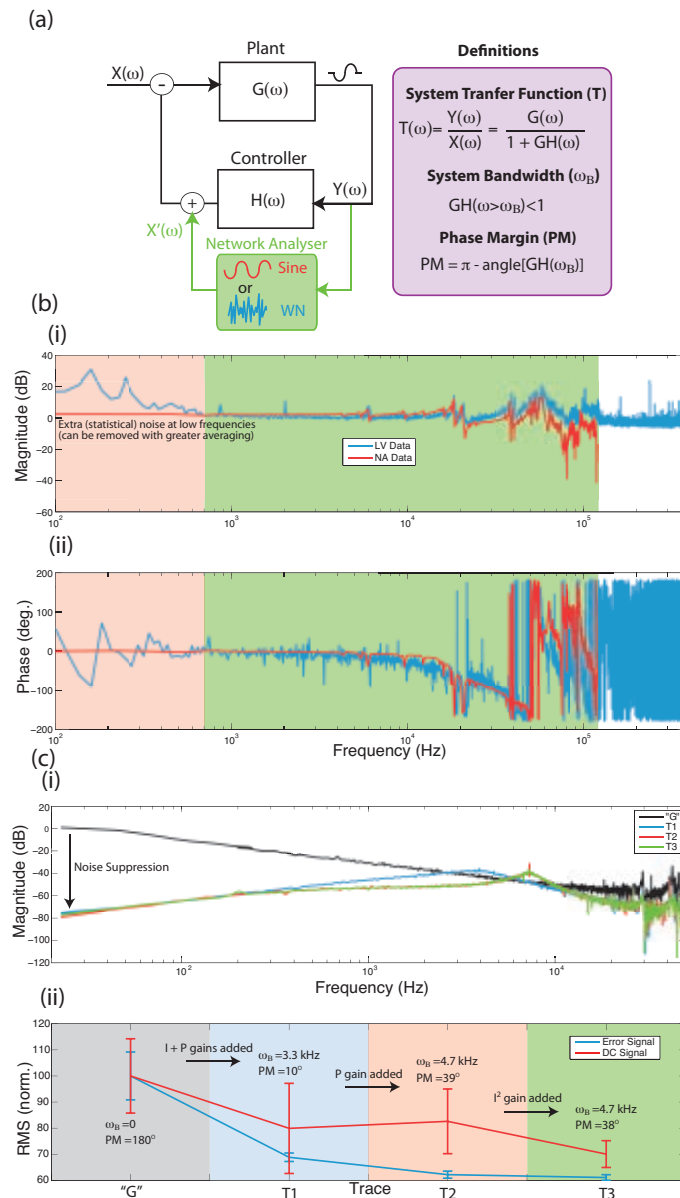


Figure B.4: Lock Optimisation. (a) Block diagram of a general closed loop feedback system with a Plant (the system to be controlled) with frequency response $G(\omega)$, and the Controller (to control the Plant using the output signal $Y(\omega)$) with frequency response $H(\omega)$. Also included is the noise added to the system $X(\omega)$. A network analyser (shown in green) can be used to measure the System Transfer Function $T(\omega)$ using either a swept sine wave (Sine) or white noise (WN), in both cases the size of the signal $X'(\omega) \gg X(\omega)$. The inset shows some of the definitions used to parametrise the locks. (b) Comparison of magnitude and phase plots for an unlocked mode-cleaner cavity using both a store bought network analyser (NA Data) and a white noise generator developed in LabVIEW® (LV Data). The traces agree very well, except for some excess noise at low frequencies for the white noise generator. (c)(i) The magnitude of T measured with different controller gains, with G being only very lightly locked to approximate the Plant and increasing gains used for T_1 to T_3 , to optimise the lock. The suppression of noise at low frequencies with the PII controller can be seen here. (ii) Root Mean Squared (RMS) measurements for both the error and DC signals for each of the traces in (i). These were taken with extra white noise present in the system and again $X'(\omega) \gg X(\omega)$. The error bars were determined from the standard deviation of the 10 values used for each set of gains. Also shown are the bandwidths and phase margins measured by the LabVIEW® code as part of its locking analysis.

the Plant transfer function $G(\omega)$ can be difficult as it requires the system to be unlocked, therefore allowing it to drift off resonance. Figure B.4(b) was taken using a Michelson interferometer set up at the rear of a cavity, not requiring the system to be locked. However, a very small amount of gain can also be used to give an approximation of G and this is what is used for the lock optimisation protocol discussed below along with H_{sim} .

B.5.2 Lock Optimisation Protocol

The idea behind the lock optimisation protocol is to find the locking parameters which allows for: large bandwidths ω_B ; large noise suppression at frequencies below ω_B ; stable locking (as measured by the phase margin); no (or minimal) excitation of resonances; and low RMS on the DC and error signals. This is achieved here mainly by investigating the transfer functions of the plant and the system (see Figure B.4). Below is a suggested method for using the code to help optimise the locking parameters if no previous G trace has been taken.

1. Turn the network analyser on using the **Take FIFO Data** switch in the **FIFO IP** cluster. Select the correct **Lock #** and make sure that **No Points in FIFO** is non-zero⁶.
2. Lock the system of interest using as little gain as possible (ideally hardly any I, and definitely no I²).
3. Use **WN On** to start sending WN to the channel of interest. For best results make sure that **WN RMS** is much larger than the system noise.
4. To display data on the **Bode Plots** tab of the **Graphs** menu make sure **Graph** is set to T and turn on the **New Data** switch in the **Data Manipulation** cluster and also click the **Start Averaging** control. There are two different averaging methods available in this code - a summation (*Sum*) which will only display new data after all traces have been collected (chosen using the **Average no. (BS)** control for 2^{BS} traces), or an exponential moving average (*EMA*) filter (with α chosen using **EMA n (Mag)**) which will constantly update but will be biased towards more recent traces. These (or no averaging) can be chosen using the **Averaging Type** control. Data will only be taken after the system have been locked for the number of counts determined by the **Lock Counter** control and displayed on the **Counter OP** indicator, this can be seen from the **Take Data** indicator. If the data on the **Bode Plots** tab vanishes, or if the **Average Counter** won't increase when **Take Data** is true, then use the **Reset Average** to start again (this problem normally occurs when the WN is switched off and then switched on again).
5. Once a trace has been created, turning **New Data** off will preserve it (if so wished). Sometimes extra gain will be present in the system, to normalise the G trace use the **T OS (dB)** control to centre the trace.
6. Once happy with the G trace, use the **G Load/Save** cluster, located in the **Saving Data** tab, to select a directory and file name to save to on the PXI hard drive. Select *Save File* in the **G Option** box and then click the **Implement G Option** switch to implement the save. This will also load the trace to the G trace bode plot

⁶For best results use the maximum number of FIFO points: 32767.

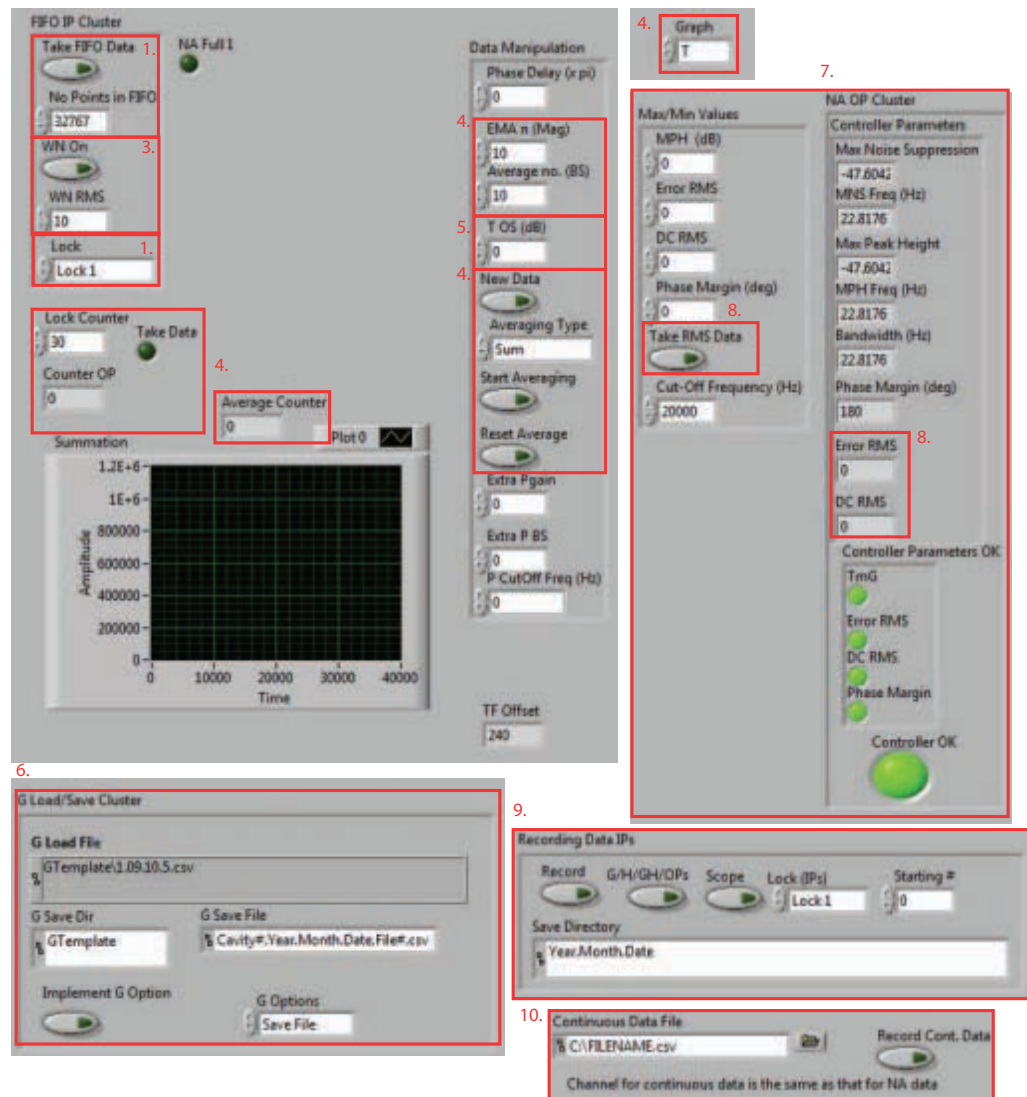


Figure B.5: Lock Optimisation Protocol. The menus and controls used to investigate and optimise a lock. The numbers and boxes refer to the instructions in text.

(which can be viewed using the **Graph** control on the **Bode Plots** tab of **Graphs** menu. Previously recorded G traces can also be loaded here.

7. Now increase the P , I , and I^2 gains to improve the lock. The tools located in the **Controller Measures** tab can be used to monitor properties of the lock such as maximum noise suppression, system bandwidth and phase margin, as well as looking at lock RMS (see next step). The **Max/Min Values** in the **Controller Measures** tab (see Section B.7.6) allow you to set limits on various parameters and the indicators let you know whether these have been achieved. The different traces can also be viewed in the **Bode Plots** tab. As with G , there may be extra gain on the controller. To take this into account, use the **Comparison Plots** tab of the **Graphs** menu to see the difference between the H simulated through Equation B.1 and H calculated through T and G (see Section B.5.1). The **Extra Pgain**, **Extra P BS**, and **P CutOff Freq (Hz)** can then be used to match the two. An example of lock optimisation is shown in Figure B.4(c). Normally this involves adding I , then P gains in an iterative process to increase noise suppression and bandwidth without exciting resonances, increasing RMS (see Section B.7.6), or lowering the PM too much, followed by I^2 gain to again increase noise suppression at low frequencies. No derivative component is included to avoid exciting resonances at higher frequencies.
8. The RMS of both the error signal and DC signal of the system can be investigated by clicking the **Take RMS Data** switch in the **Max/Min Vales** cluster of the **Controller Measures** tab. This information is then shown in the **NA OP Cluster** cluster. The RMS data is taken from the Scope, therefore to use this you must ensure that **Scope Trigger Channel** is set to the desired lock. Also, ensure that the **WN RMS** is the same when comparing RMS for different gains.
9. T , H , GH and the **Controller Parameters**, as well as scope traces and inputs for a lock can also be saved from the **Recording Data IPs** cluster in the **Saving Data** tab. Once a save directory is selected files will save in numerical order starting from 0. The saving path and next file number will be displayed in the **Recording Data OPs**. You can alter the next file number using the **Starting #** control in the **Recording Data IPs** cluster. If the file already exists the **Next File already exists** indicator will light up.
10. To investigate the long term stability of the system, the code can also take continuous data for one lock at a time in approximately 1 s time intervals. The channel that will be saved is the same channel as selected in the **FIFO IP CLuster** of the **Transfer Functions** tab.
11. To recover the saved files from the PXI hard drive you can either use the LabVIEW® Measurement and Automation Explorer (MAX) or a file transfer protocol application such as Filezilla®.

B.6 Sequential Locking

The sequential locking aspect of the code is handy if the experiment to be controlled contains many locks that depend on previous locks. The main idea behind this logic is that if any of the locks that a subsequent lock, Lock X, depends on drop lock, then Lock X will become automatically **Beam Blocked** until all the locks it depends on re-lock. This

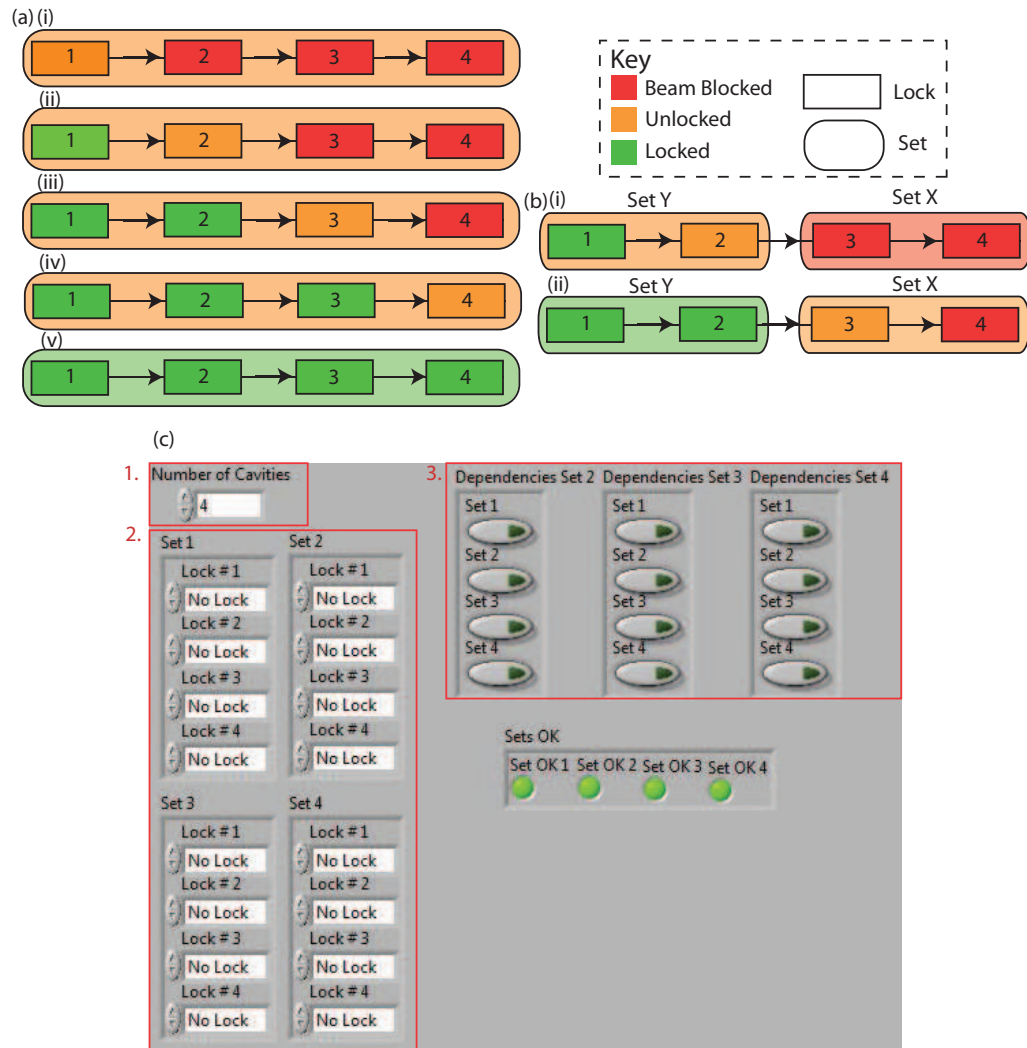


Figure B.6: Sequential Locking. A schematic diagram of the principles behind the sequential locking logic developed for the code for (a) a single set of four locks with Lock 4 depending on Lock 3 depending on Lock 2 depending on Lock 1 (as indicated by the direction of the arrows), and (b) for two sets of two locks with Set X depending on Set Y. (c) The controls used to set up the sequential locking. The numbers and boxes refer to the instructions in text.

is illustrated in Figure B.6(a). For maximum flexibility we group locks in to sets of locks. The dependence between locks in a set is defined to be linear: with Lock #4 depending on Lock #3, which depends on Lock #2, which depends on Lock #1. Sets, on the other hand, can depend on multiple other sets and all locks within a set will be **Beam Blocked** unless all locks within the sets that it depends on are locked. This is illustrated in Figure B.6 (b).

Below are the suggested steps for setting up the sequential locking for your system. All controls are situated in the **Sequential Locking** tab of the **System Inputs/Outputs** folder.

1. Set the number of locks present in the system with **Number of Cavities**
2. Set 1 is defined within the code to be the first set and therefore cannot depend on any other set. Keeping this in mind, place all locks within the provided sets. The best way of doing this is to place all those lock that follow sequentially into one set using the **Set #** clusters or, if more than four locks follow sequentially, into multiple sequential sets in the order that the locks follow on the experiment. Once you come to a lock that depends on two subsequent locks then put this lock as the first lock of a new set and continue in this fashion until all locks have been entered into a set. The sets don't have to be full, if this is the case then put *No Lock* in the remaining **Lock #X** boxes. Check the **Sets OK** indicators. If a set is not OK, make sure that one lock does not appear twice in one set, or in two different sets. Also check that you have not put more locks in sets than you selected in **Number of Cavities**. Also, *No Lock* cannot be selected as anything but the last locks in a set. Sets can span FPGA cards.
3. Now set the dependencies of the sets using the **Dependencies Set #**. If a set follows sequentially from one other set then simply select the one set, if the first lock in a set depends on two other sets then select two, etc. Check the **Sets OK** indicators are all still lit. If not make sure that there is no interdependencies (i.e. Set X depending on Set Y, which depends on Set X), and that no set depends on itself.

B.7 Menu Overview

This section provides an overview of the various menus located in the *RT_#_Lock* codes, with more in depth descriptions of the controls they contain. Most of these are located in the **System Inputs/Outputs** folder, but some are also located in the **Graphs** and **Re-ordered IP/OPs** folders.

B.7.1 Lock

Each lock menu (one for each lock present in the system) contains three input clusters and one indicator cluster. These are used to determine all parameters with respect to the individual, lock including controller gains, fmod and locking thresholds.

Signals

This cluster is for controlling the modulation/demodulation process, with:

- **BS** - the bit shift of the error signal (signal $\times 2^{BS}$) to ensure that it remains within the ± 10 V range on the scope (i.e. it does not saturate the I16 initial number range for the error signal in the slow loop);
- **Freq (MHz)** - the modulation/demodulation frequency for the signal. This value has a set step size of 0.3125 MHz to ensure that the FPGA sine look-up table and Clock. Gen. card are always at exactly the same frequency. The maximum frequency that should be used is 40 MHz, though staying below 35 MHz is advised to ensure a clean frequency output from the AD9959 card;
- **Demod Phase** - Changes the phase of the demodulating cosine (in degrees) to allow for optimisation of the error signal. The optimal phase may change between runs and so it is advised to optimise it each time the code is run;
- **Offset (V)** - due to mis-alignment etc. the error signal may not be centred around 0 V on the scope and so an artificial offset may need to be added to rectify this problem, and allow for locking to the resonance point.

Controller

The controller cluster controls the relative gains for the various parts of the controller, as well as the scan function and whether or not to lock the system, with:

- **P Gain** - gain for proportional controller;
- **P BS** - Bit Shift of P signal ($k_p = \text{PGain} \times 2^{PBS}$);
- **I Gain** - Gain of single integrator controller;
- **I BS** - Bit Shift of I signal ($K_i = \text{IGain} \times 2^{IBS}$);
- **I² Gain** - Gain of double integrator controller;
- **I² BS** - Bit shift of signal out of single integrator before entering double integrator ($k_{ii} = \text{I}^2 \text{ Gain} \times 2^{I^2BS}$);
- **Controller BS** - Common BS for P, I, and I² controllers;
- **Invert Gains** - If the signal is 180° out of phase (i.e. the high side of the error signal is on the high side of the scan, the opposite of Figure B.1(a) inset) then, instead of adding/subtracting 180° from the phase in the **Signals** cluster or using negative gains, this button can be set to true;
- **Scan Amplitude (mVpp)** - Amplitude of the sawtooth scan function in mVpp. The maximum amplitude is 20,000 mVpp, and this is only possible if **Scan Offset (mVpp) = 0** (i.e. maximum AOP range = ± 10 V);
- **Scan Offset (mVpp)** - Offset of scan from a 0 V centre point;
- **Scan Freq Step** - Determines the scan frequency in concert with **SFS Bit Shift** (see below);
- **SFS Bit Shift** - Bit shift for scan speed, with **Scan Freq (MHz)** being given by $\frac{80}{108 \times 1000} \times SFS \times 2^{SFSBS}$. See **Scan Freq (MHz)** indicator for actual scan frequency;

- **AutoLock** - If true the system will scan its full range once to set the Min and Max DC values for the Scan and Lock thresholds, before starting to lock. The system will move to Lock Mode when the DC value drops below the Scan Threshold, and will remain locked until it rises above the Scan Threshold (see below), when it will revert to Scan Mode. If the system cannot find lock after the amount of time set by **Lock TimeOut** (see Section B.7.2) then the system will undergo another complete scan to reset the maximum and minimum DC limits and try again;
- **Beam Blocked** - If the beam is going to be blocked before the cavity, set Beam Blocked to true to freeze the output signal at its current position (i.e. no change to locking signal or scan), as well as preserve the current threshold limits until it is set to false again;
- **Do Not Relock** - If you do not want to code to re-lock a lock once it has come unlocked, then select this control. When the lock comes unlocked the system will move the Scan Mode until **Do Not Relock** is deselected. Also, when this occurs, a signal is sent to the *PC_Sound_Gen.vi* code to make a sound and alert the user to this event;
- **DC Invert** - The current way the thresholds are set up is for a DC signal such as that shown in Figure B.1(a) inset, with the trough below the background signal. If the DC signal, due to detector gain/using a transmission signal etc., is inverted (i.e. a peak rather than a trough) then, by turning DC Invert to true it will invert the signal (i.e. $\times - 1$), allowing the thresholds to remain the same and **AutoLock** to function without the need to recompile the code;
- **Thresh DC(T), ES(F)** - Allows the threshold for **AutoLock** to be set from either the DC signal (T) or the absolute value of the error signals (F).

Threshold Inputs

Controls the scan and lock thresholds.

- **Scan/Lock Times** - Sets the Scan/Lock thresholds as a fraction of peak height (see next). Ideally the Scan Threshold is set to be as high as possible to avoid locking to any peaks aside from the main peak, while the lower the Lock Threshold is, the less likely the system is to move back in to Scan Mode while the system is still locked (i.e. Error Signal $\neq 0$ outside of resonance). If no DC signal is available for a lock (for instance a Michelson style lock) and/or you just want the system to lock without unlocking, then set both to 0;
- 2^{-n} - Sets n . The Scan and Lock Thresholds are set by **Scan/LockTimes** $\times 2^{-n} \times (\text{Max}-\text{Min})$ (i.e. as a fraction of the peak height). The higher they are, the closer the thresholds are to the bottom of the peak.

Threshold OP

Shows the threshold data.

- **Scan T** - The Scan Threshold limit, determined by **Scan Times** $\times 2^{-n} \times (\text{Max}-\text{Min})$. The system will enter Lock Mode when the DC value drops below the Scan Threshold;

- **Lock T** - The Lock Threshold limit, determined by **Lock Times** $\times 2^{-n} \times (\text{Max} - \text{Min})$. The system will leave Lock Mode when the DC value rises above the Lock Threshold;
- **Locked** - Tells whether the system is currently in Scan Mode (F) or Lock Mode (T), and therefore whether the system is on resonance;
- **Min** - The minimum DC value, measured in V, used to calculate thresholds (see above);
- **Max** - The maximum DC value, measured in V, used to calculate thresholds (see above);
- **DC Value** - Current DC value of the system, measured in V.

Scan Frequency (MHz)

- Displays the frequency of the scan, in MHz, calculated from **Scan Freq Step** and **SFS BS** as described above.

B.7.2 Clocks and Timing and Misc

This tab controls the frequency generator, as well the clock generator signals. Also included are the lock timeout function, output scan trigger control, and indicators relating to FPGA code timing.

Clock

Controls for Freq. Gen.

- **Resource Name** - Normally this is just the slot on the PXI chassis in which the Freq. Gen. card is located (i.e. *PXI1Slot#*). To check, use the Measurement and Automation Explorer (MAX);
- **Frequency** - This should be left at 80 MHz for clocking the ADCs. All Clock Gen. programming assumes an 80 MHz input clock;
- **Amplitude (V)** - The amplitude of the Freq. Gen. card has a small range (see LabVIEW® documentation [435]);
- **Phase** - Phase of Freq. Gen. If the input signals from the HS ADCs appear overly noisy or small, altering the clock phase may be able to fix this (see Section B.8).

Clock Board IP Cluster

Controls for all Clock Gen. boards.

- **Clk Loop Speed** - Speed at which the clock programming loop on the Master FPGA runs. This should always be greater than 107, but as the clock boards do not require fast programming, a value of 100,000 seems to work fine;
- **TTL Multiplier** - for technical reasons, the AD9959 boards used produce a much nicer output if the input 80 MHz signal is multiplied by a factor of 4. It is recommended to keep this value though it can be altered (see AD9959 documentation [434]);

- **Clk Reset** - If, for any reason, the clock generator stops working then resetting it may fix the problem. If this does not work, then the code should be stopped and rebooted. If this does not work, stop the code, turn the clock power supply off and then back on, and then start the code again (see Section B.8).

Amp/Phase Clusters

These are controls for individual clocks (as indicated). These clusters set the amplitude of the modulation signal (with a maximum of 1023), as well as the phase of the signals. It is recommended to use the **Phase** control in the **Lock** tabs to alter the phase of the error signals (see previous section). However, if the relative phase between two clock signals is important, then this is the place to alter them.

Lock TimeOut (s)

Sets the length of time (in seconds) the system will try to **AutoLock** before resetting the max and min limits and trying again.

FPGA Timing Outputs

These indicators give information about FPGA running speeds.

- **PI @ 1 MHz** - Indicates whether the FPGA code is running at the specified speed (normally 80 MHz/108 due to timing restriction of AIPs/AOPs). One indicator per FPGA;
- **PI Tick Delay** - Indicates the tick delay. Especially important if **PI @ 1 MHz = False**. One indicator per FPGA;
- **Clk @ CLS** - Indicates whether clock programming loop is able to run at specified speed (see **Clock Board IP Cluster** section above);
- **PI Clk Delay** - Indicates the clock loop tick delay. Especially important if **Clk @ CLS = False**.

Integrator Thresholds

- **I1 Max Value** - Set the maximum value for all integrator controllers in the code;
- **I2 Max Value** - Set the maximum value for all double integrator controllers in the code, if **I² Gain** is set to zero then this is ignored in the code. The actual integral values from all locks can be seen in the **Re-Ordered IP/OPs** folder in the **I Values** tab.

Scan Trig OP

Each FPGA has a scan trigger output, through AOP4, which outputs a scan signal at the same frequency as the lock selected using the **Scan Trig OP** controls, but with a set amplitude and offset. This is designed to help trigger external scopes/other pieces of hardware.

B.7.3 Sequential Locking

This tab determines the sequential locking set-up for the experiment (i.e. which locks depend on which other locks). Apart from the indicators here there are also indicators located above the **System Inputs/Outputs** folder.

- **Number of Cavities** - Input the number of cavities (or locks) present in the experiment;
- **Set #** - A set is defined here to be up to four locks which follow one after another (i.e. Lock # 4 is dependent on Lock # 3, which is dependent on Lock # 2, which is dependent on Lock # 1). This means that if Lock # 1 is not locked then locks # 2-4 will be **Beam Blocked** automatically, while if Lock # 2 is not locked then locks # 3-4 will be **Beam Blocked** etc. (see Figure B.6(a)). If less than four locks are used in a set, set the remaining spaces to *No Lock*. Sets cannot contain the same lock twice, and different sets may not contain the same lock. Sets can span across FPGA cards;
- **Dependencies Set #** - This determines which other sets Set # depends on. Note that Set 1 has no dependencies by definition. If a set is dependent on another set (i.e. Set X depends on Set Y), Set X needs all locks that make up Set Y to be locked otherwise all locks in Set X will be **Beam Blocked** (see Figure B.6(b)). Sets cannot be interrelated (i.e. Set X cannot depend on Set Y if Set Y already depends on Set X) and sets cannot depend on themselves;
- **Sets OK** - Indicates whether the sets satisfy the requirements mentioned previously in this section;
- **Set Locked** - Located above the folder, this cluster indicates whether all the locks that make up an individual set are locked;
- **BB Total** - Indicates which locks are **Beam Blocked** (either manually or automatically). These values are sent into the FPGA code;
- **All is Locked** - If all sets are locked then this will be true, otherwise it will be false.

B.7.4 Scope

To investigate the state of each lock, a scope has been included in the code. This scope (output located in the **Graphs** folder) shows the error signal, scan signal and DC signal of a lock, as well as one parameter of choice for another of the three locks on the same FPGA.

FPGA IPs

These are the scope inputs that are sent to the FPGA:

- **Trigger Direction** - Determines whether the scope triggers off the rising edge or falling edge of the scan function. The scan trigger has a set amplitude and no offset but is at the correct frequency;
- **Trigger Delay (ms)** - Determines the amount of time in milliseconds the code waits before taking data after triggering;

- **Acquisition Rate (kHz)** - Determines the rate of data collection in kHz between data points. If the scope is not triggering properly try increasing/decreasing the acquisition rate;
- **Trigger Offset (V)** - Used to move the triggering point up or down the scan function;
- **Trigger Channel** - This control is disabled. Instead, the channel the scope triggers off is determined by the control in the **Graphs** folder (see below);
- **Channel Selector** - Used to select which channels to display. Currently the code is set up so that the first three channels to be displayed are the Error Signal, Scan Trigger, and DC Signal for the lock of interest (determined by **Scope Trigger Channel**, see below). The fourth channel can be set to any other channel on the FPGA.

Other Scope Parameters

These are the scope parameters located in the RT code:

- **Scope On** - Determines whether to turn on the scope (and therefore display information on the screen) or not;
- **Scope Time Out (ms)** - Determines the time out time for the scope (if the scope is not triggered) before displaying the remaining points on the screen. If the scope screen starts to flicker, or the **Scope Full** indicator (see below) flashes, increase **Scope Time Out (ms)**;
- **RT/Offset** - Used to offset the four signals displayed on the scope;
- **RT/Scaling** - Used to scale the four signals displayed on the scope;
- **Scope Trigger Channel** - Determines which lock to trigger scope off and display channels (located in **Scope** tab in **Graphs** folder). If Lock 1 → 4 is selected then it will show data from FPGA 1, and if Lock 5 → 8 is selected then it will display data from FPGA 2;
- **Scope Channel OK** - Determines if desired trigger channel is **Beam Blocked** (if true then it will turn off the scope to avoid timing out);
- **Scope Full** - Indicates whether scope FIFO⁷ is full;
- **Scope Size** - Indicates the number of points in the scope FIFO (max of 4095).

B.7.5 Transfer Functions

This menu is to control the white noise generator and data analysis (see Section B.5.2). The data is displayed on the **Bode Plots** tab in the **Graphs** folder.

⁷first-in-first-out memory used to transfer scope data from the FPGA to the RT.

FIFO IP Cluster

These parameters control the WN generator:

- **Take FIFO Data** - Turns on Network Analyser FIFO to display information;
- **No Points in FIFO** - Number of points in Network Analyser FIFO (max of 32,767);
- **WN On** - Turns on (or off) white noise on the desired channel;
- **WN RMS** - Determines the amplitude on the WN signal;
- **Lock** - Which channel to send the white noise to, and therefore which channel to display information for. Also determines which lock information is sent to any **Continuous Data File** being recorded.

Data Manipulation

These parameters control the manipulation of data taken with the WN generator including averaging:

- **Phase Delay (x pi)** - Depending on the sign of the controller and phase of error signal, it may be necessary to add π to the current displayed T transfer function to correct the phase. This can be done here;
- **EMA n** - α for the Exponential Moving Average Filter (if used) for transfer function data ($Y_t = 2^{-\alpha} [X_t + (2^\alpha - 1) Y_{t-1}]$);
- **Average no. (BS)** - Number of traces taken for Summation average ($2^{Average_no.-(BS)}$, if used) for transfer function data;
- **T OS (dB)** - To correct for extra gain added to the current displayed T transfer function, so that 0 dB corresponds to no gain for the G trace (see sections B.5.1 and B.7.7);
- **New Data** - Lets the system take new data if set to true, otherwise will keep current data and not update when next average is complete;
- **Averaging Type** - Determines whether to use summation average (will only display new data once the number of traces indicated have been taken), or exponential moving average filter (constantly updates trace, but gives greater importance to more recent traces) to smooth transfer function data and allow for better resolution at low frequencies;
- **Start Averaging** - Must be pressed once to start system displaying data on **Bode Plots** screen;
- **Reset Average** - If system stops displaying data (normally when white noise is turned off and then back on again), hit reset to start displaying data again;
- **Extra Pgain** - As was the case with extra T OS, sometimes there seems to be a constant offset between the calculated H value (from T and G) and the simulated H value (using Equation B.1). The different between the two can be seen from the **Comparison Plots** tab of the **Graphs** folder. The **Extra Pgain** can be used to compensate for this;

- **Extra P BS** - As was the case with the controller parameters (see Section B.7.1), this BS is used in tandem with the **Extra Pgain**;
- **P CutOff Freq** - Determines the frequency range (starting from 0 Hz to **P CutOff Freq**) for the compensation.

Other Transfer Function Parameters

- **Graph** - determines which parameter to plot on **Bode Plots** tab of **Graphs** folder, where it is located. Can display T , G (if taken and saved/loaded, see Section B.5.2), H , GH , as well as current error signal (top graph) and white noise (bottom graph) traces;
- **Lock Counter** - To ensure that data is only taken when the system is locked and stable, the **Lock Counter** enables users to set a number of counts the system must have been locked for before data will be taken. If this is not an issue, simply set to 0;
- **Counter OP** - Indicates current # of **Lock Counter** (up to Max);
- **Take Data** - Indicates whether **Lock Counter** currently enables data to be taken;
- **Average Counter** - Displays the number of averages that have been taken so far in *Sum* mode;
- **Summation Graph** - Displays current Summation average data. If this goes blank in *Sum* mode then it is necessary to use **Reset Average**;
- **NA Full** - Indicates whether network analyser FIFO is full.

B.7.6 Controller Measures

This menu is for displaying analysis for the lock of interest (see Section B.5.2).

Max/Min Values

This is where the user can set their desired max/min values for certain locking parameters:

- **MPH (dB)** - Input the maximum peak height which can be tolerated for the lock up to **Cut-Off Frequency (Hz)** (see below);
- **Error RMS** - Input the maximum error signal RMS that can be tolerated;
- **DC RMS** - Input the maximum DC RMS that can be tolerated;
- **Phase Margin (deg)** - Input the minimum phase margin that can be tolerated up to **Cut-Off Frequency (Hz)** (see below);
- **Take RMS Data** - Click to take an RMS data set for both DC and error signals of the lock currently displayed on the scope;
- **Cut-Off Frequency (Hz)** - Input a frequency that will be a maximum for the locking analysis values (i.e. system will ignore all frequencies above this).

NA OP Cluster

This cluster displays the locking parameters measured by the system:

- **Max Noise Suppression** - Displays the value of the maximum noise suppression of the system (below **Cut-Off Frequency (Hz)**);
- **MNS Freq (Hz)** - The frequency at which the maximum noise suppression occurs;
- **Max Peak Height** - Displays the value of the maximum peak height of the system (below **Cut-Off Frequency (Hz)**);
- **MPH Freq** - The frequency at which the maximum peak occurs;
- **Bandwidth (Hz)** - Displays the bandwidth of the system (i.e. the unity gain frequency ω_B of the GH graph);
- **Phase Margin (deg)** - Displays the phase margin of the system (i.e. the phase of GH above $-\pi$ at ω_B);
- **Error RMS** - Displays the Error signal RMS value from the previous RMS run;
- **DC RMS** - Displays the DC signal RMS value from the previous RMS run;
- **Controller Parameters OK** - Indicators to show whether the system parameters agree with the **Max/Min Values** entered (**TmG** = max peak height of T);
- **Controller OK** - Indicates whether all four of the above indicators agree or not with the **Max/Min Values** entered.

B.7.7 Saving Data

This menu allows for saving of data from different sources - either one off or continuous.

Recording Data IPs/OPs

This cluster is for recording transfer function data:

- **Record** - Press to save current T trace (as .csv), along with f - the frequency axis scale;
- **G/H/GH/OPs** - Turn on to also save these traces along with T . Outputs are those located in the **Controller Parameters** cluster;
- **Scope** - Will also save scope data;
- **Lock IPs** - Will also save a subset of the lock inputs if selected (In order: **Signal/BS**; **P Gain**; **I Gain**; **I² Gain**; **I² BS**; **Controller BS**; **P BS**; and **I BS**);
- **Save Directory** - Determines which directory to save data to (in RT home directory on the PXI chassis hard drive). You should always create a new directory for each set of data (i.e. # 0-X) to avoid saving over previous data, see below;
- **Starting #** - Determine starting number for traces (will increase by 1 after each data is saved);

- **appended path** - Displays the path of the previous saved file, will be of form *SaveDirectory/File#/T.csv* etc.);
- **Next File #** - Displays the number of the next file # to save (can be changed by changing **Starting #**);
- **Next File already exists** - This indicator tell you if the next file # already exists on the PXI hard drive. If you then proceed to take data the file will be saved over.

Continuous Data

This cluster is for saving continuous data, in $\simeq 1$ s intervals. It will record the DC value of the lock selected for **FIFO IP Cluster** in the **Transfer Functions** tab, as well as the time interval between points in ms. It will only save data if **Take Data** is true (see Section B.7.5). If this is not desirable then set **Lock Counter** to 0.

- **Continuous Data File** - To name and select type (by changing extension, i.e. .csv) of file to save;
- **Record Continuous Data** - Turn on to take data and turn off to stop taking data;
- **Next Cont. File Exists** - This indicator tells you if your selected file for the next set of continuous data already exists on the PXI hard drive. If you proceed to take data then it will be saved over the top of the previous data.

G Load/Save Cluster

This cluster is for saving *G* data to allow for locking parameters to be calculated (see Section B.5.2).

- **G Save Dir** - Determines directory (inside RT home directory) to save *G* data to;
- **G Save File** - Determines the file name (and extension, i.e. .csv) for *G* save file;
- **G Load File** - Determines the directory and file name (as well as extension) for *G* file to load;
- **G Options** - Determine desired *G* operation (save/load file);
- **Implement G Option** - Push to implement currently selected **G option**.

B.7.8 Re-ordered IP/OPs and Errors

The folder on the right hand side of the **Graphs** folder is where all variables are placed before they are re-bundled into a user friendly format, and they are only included here as they must be included somewhere for coding reasons and for initial debugging. The only tab which provides useful information is the **Errors** tab. Above the **System Inputs/Outputs** folder is an **Error** indicator. If this flashes true then it means that an error has been detected for either the Freq. Gen., one of the FPGAs (most probably to do with the FIFO read out and time-out), or with the continuous data saving code. To determine which, look at the **Errors** tab and one should display an error code and message to help fix the problem.

B.7.9 PC_Sound_Gen.vi

This code, located in the PC component of the project, can be used in conjunction with the RT locking codes to produce sounds due to various events. The RT locking codes communicate to the *PC_Sound_Gen.vi* code using Shared Variables. Currently the only event that will register is if a lock becomes unlocked when **Do Not Relock** is selected, but this can be altered in the RT code. To make use of the sound generation option simply run this code in conjunction with the RT locking code.

Gen Sound IP Cluster

These values have been preset and it is inadvisable to change any of these controls, except the **Frequency** input, to alter the pitch of the signal.

Sound Format

Again, these values have been preset and it is inadvisable to change any of these unless there is no sound being produced or errors occur.

Other Controls and Indicators

- **Volume** - Changes the volume of the sound, ranging from 0-100;
- **Test** - Toggle to test out speakers (sound should be generated when **Test** = true);
- **Mute** - To stop sound when code is running put **Mute** = true;
- **Length** - Sets length of sound signal: the greater the value the longer each sound will be and the slower the refresh rate;
- **Came Unlocked** - This indicates when a sound event (see above for current sound event) occurs in the RT.

B.8 Known Problems

Though we have done our best to remove any bugs from the code, there are still some issues that we have not fixed. These are listed below.

- **Extra Noise on Error Signals** - This was mentioned in Section B.7.2. Though the Master FPGA code is set up to trigger the Freq. Gen., sometimes the signal from the Freq. Gen. does not trigger the HS ADCs in sync with the running of the FPGA code, leading to a signal with excess noise. This can be removed by scanning the phase on the Freq. Gen., though if multiple HS ADCs are being run with the same source and the cable lengths are not the same it is possible that there will not exist a region where all HS ADCs are in sync. This is something that should be considered when setting up the HS ADCs.
- **Real Time Time-Out** - Sometimes there is an unexplained loss of connection between the PC and the Real-Time controller (normally with a message such as: *Waiting for Real-Time target (rt PXI target) to respond*). The code continues to run on the FPGAs and the RT when this happens, it just does not allow for updates to/from the PC. Connection can be re-established by re-deploying the RT code.

We have found this a major issue with LabVIEW® 2009, but seems to be fixed in LabVIEW® 2010.

- **Clock Gen. Issues** - More of a hardware problem, however, we have found that sometimes the AD9959 DDS boards will stop sending out their modulation signals. This normally occurs when there is activity around their power supply and we hope that moving to a shielded rack mountable system will fix this problem. In the mean time, the best way to fix the problem is to stop the code, turn off and then back on the clock power supply, and then re-deploy the code.
- **AutoLock Not Working** - This can occur for a few reasons: The first step would be to check that the phase of the error signal is correct (and not 180° out), after this you may want to consider lowering both the Scan Threshold and Lock Threshold (this is especially the case for small DC signals on a large background) to allow for a greater locking region. Sometimes also the scan is too fast to allow for the controller to take effect before the Lock Threshold is reached and you may want to slow this down. If the system starts locking to the top or bottom of the error signal rather than the zero point this indicates that you are using too much gain and that you should play with the ratios of the three different gains to try to rectify this.
- **Scope Not Triggering Correctly** - The first thing to check is that the selected scope channel is not beam blocked. If the scope is displaying a signal but not triggering, try increasing/decreasing the acquisition rate. If the scope is flashing, try increasing the scope time out or increasing the scan frequency.
- **Code Not Compiling** - The FPGA code for *PHD_Lock_5001* is very close to filling the FPGA memory, therefore if you are trying to add extra logic to the code then there is every likelihood that it will not compile the first or second times. Try at least three times to make sure. If that doesn't work then a LabVIEW® 2009 version of the FPGA code is available upon request, as we have found that the code compiles much more easily on the earlier version (but runs much better on 2010).

B.9 Acknowledgements

I would like to thank all my fellow authors on Ref. [433]: Helen Chrzanowski; David Parrain; Dr. Ben Buchler; Prof. Ping Koy Lam; and Dr. Thomas Symul, for all their help and support with this project. I would also like to thank Dr. Dan Shaddock for helpful hints given, and Neil Hinchey and Oliver Thearle for their help developing the hardware to accompany the code.

Cavities and Pound-Drever-Hall Locking

Cardinal Ximénez: “NOBODY expects the Spanish inquisition!”

- “The Spanish Inquisition” *Monty Python’s Flying Circus*, BBC, 22 September 1970, TV Series.

In this appendix we will briefly go through the theory behind producing an anti-symmetric error signal to frequency stabilise a cavity (as demonstrated in Ch. 4). The mathematical formulation below is based on those in Ref.s [206] and [436].

C.1 Cavities

A cavity is a resonator for laser light, in which the light will perform a number of cycles. The simplest case is a cavity created by two mirrors, both with reflectivity r and transmissivity t , however here we will start with a three-mirror cavity (shown in Fig. C.1) as this is the main type of cavity used throughout this thesis. If the back mirror M_3 is assumed to have a reflectivity of 1 (i.e., totally reflective) then the two are equivalent. The main properties of a cavity are defined below

$$F_c = \frac{4r^2}{(1-r^2)^2} \quad (\text{C.1})$$

$$\mathcal{F} = \frac{\pi\sqrt{F_c}}{2} \quad (\text{C.2})$$

$$\Delta f_{FSR} = \frac{2}{2nL} \quad (\text{C.3})$$

$$Q = \frac{2L}{\lambda}\mathcal{F}, \quad (\text{C.4})$$

where L is the length of the cavity round trip, n is the refractive index of the cavity medium (taken here to be air, therefore $n = 1$) and λ is the wavelength of the light in the cavity. These terms are: the coefficient of finesse, finesse, free spectral range (FSR), and Q-factor respectively. The finesse tells us about the ratio of transmission peak size Γ to a full period by $\Gamma = 2\pi/\mathcal{F}$. The FSR tells us how far we can scan in frequency before coming across the same mode again. While the last of these determines how many round trips a photon will undergo inside a cavity, as mentioned in Ch. 1 when discussing cavities as potential quantum memories.

If we consider light entering the cavity in Fig. C.1(a) we see that, of the input light field \mathcal{E}_{in} , an amount $r\mathcal{E}_{in}$ is immediately reflected and an amount $\mathcal{E}_o = t\mathcal{E}_{in}$ enters the

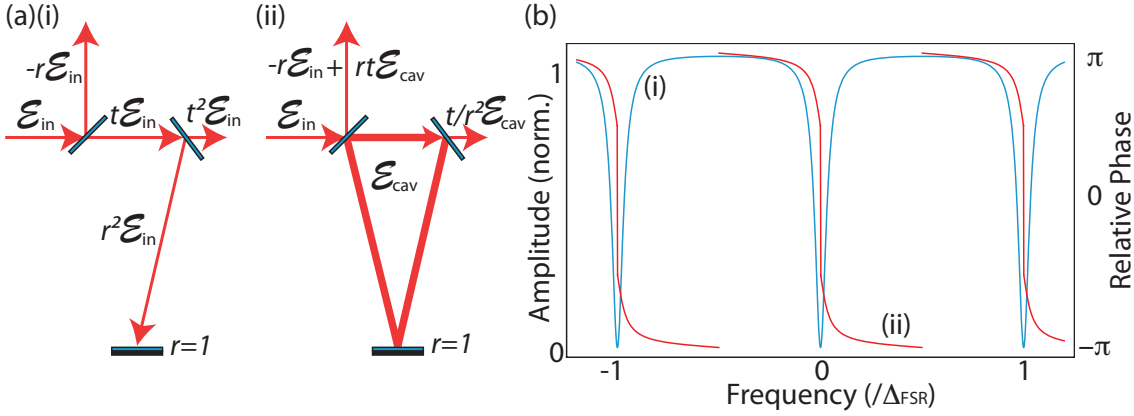


Figure C.1: Cavity Theory. (a) Build up of optical power inside the cavity from (i) to (ii) for on-resonance light. (b) Plot of (i) Power (blue, left-hand scale) and (ii) phase (red, right-hand scale) for the reflected signal from Eq. C.9 with $r = 0.9$.

cavity. Some of this light will pass straight out of the cavity, while an amount $r^2\mathcal{E}_{in}$ will circle around and join with the next input to the cavity (assuming no loss). Also, if the light is not resonant with the cavity (i.e., $L \neq N\lambda$, where N is an integer), then it will experience a phase shift with respect to the input field of $\phi(\omega) = 2\omega L/c$. Therefore, after the first round trip the signal will be

$$\mathcal{E}_1 = r^2 \text{Exp}[-i\phi(\omega)] \mathcal{E}_0. \quad (\text{C.5})$$

This relation will hold between any two adjacent fields j and $j + 1$ and therefore the total field inside the cavity will be given by

$$\mathcal{E}_{cav}(\omega) = \sum_j^{\infty} \mathcal{E}_j = \frac{\mathcal{E}_0}{1 - r^2 \text{Exp}[-i\phi(\omega)]}. \quad (\text{C.6})$$

The total power in the cavity is given by

$$\begin{aligned} P_{cav}(\omega) &\propto |\mathcal{E}_{cav}(\omega)|^2 \\ &= \frac{|\mathcal{E}_0|^2}{(1 - r^2)^2 + 4r^2 \text{Sin}[\phi(\omega)/2]} \\ &\leq \frac{|\mathcal{E}_0|^2}{(1 - r^2)^2}, \end{aligned} \quad (\text{C.7})$$

the maximum value for P_{cav} occurring on resonance (i.e., $\phi(\omega) = 2\pi N$, N again being an integer). We can also calculate the transmitted and reflected fields from \mathcal{E}_{cav}

$$\mathcal{E}_{trans}(\omega) = \frac{t \text{Exp}[-\phi(\omega)]}{r^2} \mathcal{E}_{cav} \quad (\text{C.8})$$

$$\begin{aligned} \mathcal{E}_{ref}(\omega) &= -r\mathcal{E}_{in} + rt\mathcal{E}_{cav} \text{Exp}[-\phi(\omega)] \\ &= \mathcal{E}_{in} \left(\frac{-r(1 - \text{Exp}[-i\phi(\omega)])}{1 - r^2 \text{Exp}[-i\phi(\omega)]} \right) \\ &= \mathcal{E}_{in} R(\omega), \end{aligned} \quad (\text{C.9})$$




Figure C.2: PDH Error Signals (a) Error signal for low modulation frequency, using $\text{Re}[\alpha(\omega, \omega_{mod})]$ with $\omega_{mod} = .001\Delta_{FSR}$. (b) Error signal for high modulation frequency, using $\text{Im}[\alpha(\omega, \omega_{mod})]$ with $\omega_{mod} = .05\Delta_{FSR}$.

where we define $R(\omega)$ as the reflection coefficient. The power and phase of the reflected component of the light field are plotted in Fig. C.1(b). As can be seen, the normalised reflected signal is at one except around resonance (while the transmitted signal is at zero except around resonance) and the phase is zero at resonance and anti-symmetric around resonance.

C.2 Pound-Drever-Hall Locking

Unfortunately, from a locking point of view, both the reflected and transmitted signals are symmetric around resonance. They are therefore not useful as an error signal. One option would be to lock the cavity half-way up a peak but this introduces issues such as the zero-crossing for the error signal being dependent on the laser power, not to mention the fact that we will not be locking to resonance. This is where the Pound-Drever-Hall locking protocol [246, 247] comes into play. The first step is to modulate the phase of the light. This is usually achieved using an electro-optic modulator (EOM) to give

$$\mathcal{E}_{mod}(\omega, \omega_{mod}) = \mathcal{E}_{in} \text{Exp}[-i(\omega t + \beta \text{Sin}[\omega_{mod} t])] \quad (\text{C.10})$$

$$\begin{aligned} &= \mathcal{E}_{in} (J_0(\beta) \text{Exp}[-i\omega t] \\ &\quad + J_1(\beta) \text{Exp}[-i(\omega + \omega_{mod}) t] \\ &\quad - J_1(\omega - \omega_{mod}) \text{Exp}[-i(\omega - \omega_{mod}) t]), \end{aligned} \quad (\text{C.11})$$

where J_k is a Bessel function of order k . Combining the above equation with Eq. C.9 gives us

$$\begin{aligned} \mathcal{E}_{ref}(\omega, \omega_{mod}) &= \mathcal{E}_{in} (R(\omega) J_0(\beta) \text{Exp}[-i\omega t] \\ &\quad + R(\omega + \omega_{mod}) J_1(\beta) \text{Exp}[-i(\omega + \omega_{mod}) t] \\ &\quad - R(\omega - \omega_{mod}) J_1(\omega - \omega_{mod}) \text{Exp}[-i(\omega - \omega_{mod}) t]). \end{aligned} \quad (\text{C.12})$$

We can then calculate the power reflected $\mathcal{E}_{ref}\mathcal{E}_{ref}^*$ to be

$$\begin{aligned}
P_{ref} &= P_c |R(\omega)|^2 + P_{sb} \left(|F(\omega + \omega_{mod})|^2 + |F(\omega - \omega_{mod})|^2 \right) \\
&\quad + 2\sqrt{P_c P_{sb}} \left(\text{Re} \{ R(\omega) R^*(\omega + \omega_{mod}) - R^*(\omega) R(\omega - \omega_{mod}) \} \text{Cos} [\omega_{mod}t] \right. \\
&\quad \left. + \text{Im} \{ R(\omega) R^*(\omega + \omega_{mod}) - R^*(\omega) R(\omega - \omega_{mod}) \} \text{Sin} [\omega_{mod}t] \right) \\
&\quad + O(2\omega_{mod}), \tag{C.13}
\end{aligned}$$

where $P_c = J_0^2(\beta) |\mathcal{E}_{in}|^2$ is the power in the carrier, $P_{sb} = J_1^2(\beta) |\mathcal{E}_{in}|^2$ is the power in the sidebands, and $O(2\omega_{mod})$ contains all terms with frequency terms greater or equal to $2\omega_{mod}$. The above equation simplifies if we take the limit $\omega_{mod} \ll \Delta_{FSR}/\mathcal{F}$, where only the real part of $R(\omega)R^*(\omega + \omega_{mod}) - R^*(\omega)R(\omega - \omega_{mod}) = \text{Re}[\alpha(\omega, \omega_{mod})]$ exists. This is plotted in Fig. C.2(a). As can be seen, this is anti-symmetric around resonance and can therefore be used as our error signal. To extract only this term from Eq. C.13 we multiply it by $\text{Cos} [\omega_{mod}t + \theta]$ to give

$$\begin{aligned}
P_{ref} &= \mathcal{K} \text{Cos} [\omega_{mod}t + \theta] + 2\sqrt{P_c P_{sb}} \text{Re} [\alpha(\omega, \omega_{mod})] \text{Cos} [\omega_{mod}t] \text{Cos} [\omega_{mod}t + \theta] \\
&\quad + O(2\omega_{mod}) \\
&= \mathcal{K} \text{Cos} [\omega_{mod}t + \theta] - \frac{\text{Re} [\alpha(\omega, \omega_{mod})] (\omega, \omega_{mod})}{2} \text{Cos} [\theta] \\
&\quad + \frac{\text{Re} [\alpha(\omega, \omega_{mod})] (\omega, \omega_{mod})}{2} (\text{Cos} [2\omega_{mod}t] \text{Cos} [\theta] \\
&\quad - \text{Sin} [2\omega_{mod}t] \text{Sin} [\theta]) + O(2\omega_{mod}), \tag{C.14}
\end{aligned}$$

where we have used \mathcal{K} to represent the constant terms. If we were to then perform a low-pass filtering operation on the above equation to remove all non-dc terms, we obtain

$$-\frac{\alpha(\omega, \omega_{mod})}{2} \text{Cos} [\theta]. \tag{C.15}$$

This gives us our anti-symmetric error signal, optimised at $\theta = 0, \pi$.

At high modulation frequencies $\omega_{mod} \gg \Delta_{FSR}/\mathcal{F}$ we find that only the imaginary component of $\alpha(\omega, \omega_{mod})$ exists. This is again anti-symmetric, and is plotted in Fig. C.2(b). We can again isolate this term by multiplying with $\text{Cos} [\omega_{mod}t + \theta]$, except now the error signal will be optimised at $\theta = \pi/2, 3\pi/2$.

In Ch. 4 we discuss how we can use these anti-symmetric error signals to provide feedback control to the cavities.

Mathematica Codes

Mr. Finch: "I designed the machine to see acts of terror, but it sees everything..."
 - *Person of Interest* (opening titles), CBS, 2011-Present, TV Series.

In this appendix we present some of the Mathematica codes used for calculations throughout this thesis. Firstly an example of the code used to determine the currents for the multi-element coil from Ch. 7, followed by example of the code used to extract information from heterodyne traces used in Ch.s 7 and 12. Finally we present example code used for the ac Stark shift calculations presented in Ch. 10.

D.1 Code to Calculate Multi-Element Coil Currents

... for a two-part linear gradient.

```
(*Running instructions: enter values below, quite Kernal and then run all*)
(*Input your desired magnetic field bandwidth across cell (DB)
- put -ve for opposite grad - and cell length (length)*)

DB1 = -1500; (*initial gradient, and change in B field in 1st part of cell in mG*)
DB2 = 500; (*change in B field in 2nd part of cell in mG*)
flength = 200; (*length in mm*)
switchpoint = 100; (*point @ where gradients change over in mm*)
offsetX = 0; (*offset of field from 0 mm*)
offsetB = 750; (*initial B field offset in G @ centre of cell in mG*)
stepsize = 1; (*step size for table in mm*)

(*Required Constants*)
kb = 1.3806503 *10^-23 (*m2 kg s-2 K-1*);
m = 0.085/(6.022*10^23) (* Kg *);
\[CapitalGamma] = 38.11 10^6 (* s-1 *);
\[HBar] = 1.05457148 10^-34 (*m2 kg/s*);
kl = 2 \[Pi] 1281654.938993 (* m-1*);
\[Mu]b = 9.274008 10^-24 (* J/T *);
\[Mu] = 4 \[Pi] 10^-7 (*H/m*);
Is = 1.64 (* mW/cm^2 *);
n = 8; (*Number of coils*)

(*Initial Coil Conditions*)
```

```

r = 0.042; (*cell radius in m*)
helicity = 1;
\[\Delta] = .002; (*spacing between coils*)
Tlength = 0.28;    (*Total length of the coil array*)

length = (Tlength - \[\Delta] (n - 1))/n ;  (*"length of each coil"*)
ci = length / (2 \[Pi]);    (* length = 2\[Pi]c *)
overlap1 = 0.0;
(*overlap /overhang of the first and last coil past the range of the normal length *)
overlap2 = 0.0;

(*Calculations*)
DBT1 = DB1*10^(-7); (*DB in Tesla*)
DBT2 = DB2*10^(-7);
cellLength = flength*10^(-3); (*length in m*)
switchpointM = switchpoint/1000; (*switch point length in m*)
slope1 = DBT1/cellLength; (*slope in T/m*)
slope2 = DBT2/cellLength; (*2nd slope in T/m*)
intercept1 = -DBT1/2 + offsetB*10^(-7); (*y intercept*)

line1[x_] = slope1*x + intercept1; (*desired linear gradient*)
intercept2 = line1[switchpointM] - slope2*switchpointM; (*y intercept2*)

line2[x_] = slope2*x + intercept2;
x0 = offsetX*10^(-3);
y0 = x0;
y1 = x0 + switchpointM;
xf = x0 + cellLength;
y2 = xf;
line[x_] := \[Piecewise]
  { {slope1*x + intercept1, y0 <= x <= y1}, {slope2*x + intercept2, y1 < x <= y2} }
line[0]
line[0.1]
line[0.2]
stepSize = stepsize*10^(-3);
DesiredField = Table[{x, line[x]}, {x, x0, xf, stepSize}];
Hit = ListPlot[DesiredField];

<< VectorAnalysis`
Clear[a, b, c, d, A, \[Beta], \[Gamma], Cur, CurS, \[Chi], Current, p, z1];
A = Array[a, n];(* A = {a[1],a[2],a[3],a[4],a[5],a[6],a[7],a[8]} *)
\[Beta] = Array[b, n];
\[Gamma] = Array[d, n];
\[Chi] = Array[c, n];

Cur = {};
For[i = 1, i <= n, i++,
  Cur = Append[Cur, Symbol[StringJoin["Cur", ToString[i]]]];

```

```

B[z1_, Cur_, a_, b_, d_, c_] := 1/(10000000 c^2) Cur helicity ;
(* Total Field is the sum of the contributions from each coil at each \
z *)
Bt[z_, Cur_, A_, \[Beta]_, \[Gamma]_, \[Chi]_] =
  Sum[B[z, Symbol[StringJoin["Cur", ToString[i]]], a[i], b[i], d[i],
    c[i]], {i, n}];
InitialParameters = {};
For[i = 2, i < n, i++,
  InitialParameters = Join[InitialParameters, {a[i] -> 0, b[i] -> 4,
    d[i] -> (i - 1) (length + \[Delta] ) - (Tlength - cellLength)/2,
    c[i] -> ci}]];
InitialParameters = Join[InitialParameters, {c[1] -> ci + overlap1/(2 \[Pi]) ,
  a[1] -> 0, d[1] -> -overlap1 - (Tlength - cellLength)/2, b[1] -> 4}];
InitialParameters = Join[InitialParameters, {c[n] -> ci + overlap1/(2 \[Pi]) ,
  a[n] -> 0, d[n] -> (i - 1) (length + \[Delta] ) - (Tlength - cellLength)/2,
  b[n] -> 4}];
InitialParameters
(*d[1] = -overlap1; *) (* position of first coil*)

(* d[i] are Initial positions of the coils : d[i] is the leftmost position
of the i'th coil *)

(*Plots of the desired fields*)
Hit = ListPlot[DesiredField];
cons = { -4 < Cur1 < 4 , -4 < Current1 < 4 , -4 < CurS1 < 4 };
fit = NonlinearModelFit[
  DesiredField , {Bt[z, Cur, A, \[Beta], \[Gamma], \[Chi]] /.
    InitialParameters}, Cur, z]

params = fit["BestFitParameters"]

"Error Analysis"
Errors = fit["FitResiduals"][[ Round[-x0/stepSize + 1] ;;
  Round[(-x0 + cellLength) /stepSize + 1]]];
Norm [Errors]^2
fit["EstimatedVariance"]
fit["ANOVATable"]
Miss = Plot [fit[x], {x, -.1, .3}, PlotStyle -> {{Thickness[0.005], Hue[0]}}];
hitandmiss = Show[{Miss, Hit}, PlotRange -> All]
(*for saving to file*)
CA = 1000*{params[[1, 2]], params[[2, 2]], params[[3, 2]], params[[4, 2]],
  params[[5, 2]], params[[6, 2]], params[[7, 2]], params[[8, 2]]};

(*Saving Files with Consistent Names*)
FileName = "DLin_" <> ToString[DB1] <> "mG_" <> ToString[DB2] "mG_" <>
  ToString[offsetB] "mGOS_" <> ToString[switchpoint] <> "mmSP";
Export["/Users/bensparkes/Documents/PhD/Data, Analysis and Simulations
/10.04 - Frequency Filtering/11.05 - Mathematica Notebooks/"

```

```
<>ToString[FileName]<>".csv",CA]
```

D.2 Heterodyne Data Extraction Code

The first section of code is for loading traces saved in .txt form from a LeCroy oscilloscope to remove text. If data is just two columns, with the first containing the time information and the second containing the heterodyne data, then some of this step is not needed (files still have to be named correctly).

```
(*Loading multiple files for each runfor resaving in .csv format*)

d = 1; (*setting which file to take*)
dmax = 14; (*number of runs*)
(*defining null arrays*)
fileext = Table[0 k, {k, 1, dmax}];
FileNumbers = Table[0 k, {k, 1, dmax}];
date = Table[0 k, {k, 1, dmax}];
(*Entering files names - saved with date, run number and file number*)
date[[1]] = 121114;
fileext[[1]] = "01nm"; (*file name*)
FileNumbers[[1]] = Table[i, {i, 1, 10}]; (*which/how many files to load*)
(*etc*)
channel = C1;
numberoffiles = 10; (*d=4 has only 19, 10 for all else*)

(*For creating names for all files - only needs to be done once*)
SetDirectory["DIRECTORY/"];
Directory[];
ghgh[j_] = If[j > 99, "00", If[j > 9, "000", "0000"]];
filenames =
  Table[ToString[channel] <> ToString[date[[d]]] <> "-" <>
    ToString[fileext[[d]]] <> ToString[ghgh[FileNumbers[[d]][[i]]]] <>
    ToString[FileNumbers[[d]][[i]]] <> ".txt", {i, 1, numberoffiles}];
filenames2 =
  Table[ToString[channel] <> ToString[date[[d]]] <> "-" <>
    ToString[fileext[[d]]] <> "000" <>
    ToString[FileNumbers[[d]][[i]]] <> "a", {i, 1, numberoffiles}];

(*For saving new files, comment out for reloaded files*)
Do[
  file1 = Import[ToString[filenames[[a]]]];
  Export["/Users/bensparkes/Documents/PhD/Data, Analysis and
  Simulations/12.02 - Real Cold Atom GEM/Scope (HP) Traces/" <>
    ToString[filenames2[[a]]] <> ".csv", file1];
  file1c = Import[ToString[filenames2[[a]]] <> ".csv"];
  jmin = 6;
  jmax = Dimensions[file1c][[1]] - 1;
  lmax = jmax - 5;
```

```

file2 = Table[file1c[[j, k]], {j, jmin, jmax}, {k, 1, 2}];
Export["/DIRECTORY" <>
  ToString[filenames2[[a]]] <> ".csv", file2]
, {a, 1, numberoffiles}]

(*Loading Data for analysis*)
Clear[b, amp, jmin, \
jmax, lmax, data]

data1 = Table[0 i, {i, 1, numberoffiles}];

SetDirectory["/DIRECTORY"];

Do[
  file1b = Import[ToString[filenames2[[b]]] <> ".csv"];
  jmin = 1;
  jmax = Dimensions[file1b][[1]];
  file2 = Table[file1b[[j, 2]], {j, jmin, jmax}];
  data1[[b]] = file2;
  , {b, 1, numberoffiles}]

(*Plotting data to find out bounds*)
b = 1;
bnc = Table[0 k, {k, 1, 23}];
enc = Table[0 k, {k, 1, 23}];
bnc[[1]] = 5000;
enc[[1]] = 25000;
(*etc*)
Show[ListLinePlot[data1[[b]], PlotRange -> All]]
Show[ListLinePlot[data1[[b]][[bnc[[d]] ;; enc[[d]]]],
  PlotRange -> All]]

(*Removing Mean & shortening and gestimating as
many params as possible out*)
Clear[b]
amin = 1;
amax = numberoffiles
data = Table[0 i, {i, amin, amax}];
datashort = Table[0 i, {i, amin, amax}];
amp = Table[0 i, {i, amin, amax}];
centrepos = Table[0 i, {i, amin, amax}];
\[Epsilon] = Table[0 i, {i, amin, amax}];
tmax = Table[0 i, {i, amin, amax}];

Do[
  datamean = Mean[data1[[b]][[bnc[[d]] ;; enc[[d]]]]];
  data[[b]] = data1[[b]] - datamean;
  datashort[[b]] = data[[b]][[bnc[[d]] ;; enc[[d]]]];

```

```

tmax[[b]] = Dimensions[datashort[[b]]][[1]];
amp[[b]] = Max[Abs[datashort[[b]]]];
centrepos[[b]] = Position[Abs[datashort[[b]]], amp[[b]][[1, 1]];
\[Epsilon] [[b]] = If[datashort[[b]][[centrepos[[b]]]] > 0, 1, -1];
, {b, 1, numberoffiles}]

(*fitting*)
(*Best Guesses*)
sigma = Table[0 i + 4000, {i, 1, dmax}];
freq = Table[0 i + 0.001, {i, 1, dmax}];

\[Theta] = 0;
(*creating null tables*)
Clear[b, i]
toF = Table[0 i, {i, amin, amax}];
sigmaF = Table[0 i, {i, amin, amax}];
\[Theta]F = Table[0 i, {i, amin, amax}];
freqF = Table[0 i, {i, amin, amax}];
ampF = Table[0 i, {i, amin, amax}];
OSF = Table[0 i, {i, amin, amax}];
fittable = Table[0 i, {i, amin, amax}];
errorfit = Table[0 i, {i, amin, amax}];

Do[
  sigmaF[[b]] = sigma[[d]];
  toF[[b]] = centrepos[[b]];
  freqF[[b]] = freq[[d]];(*for fitting single gaussian*)

  ampF[[b]] = amp[[b]]; (*for if cannot fit!*)
  (*fitting function*)
  Fittingnow =
  NonlinearModelFit[datashort[[b]],
    ampf Exp[-0.5 (t - tof)^2/(sigmaf^2)] \[Epsilon] [[b]] Cos[
      2 \[Pi] freqf (t - thetaf)] +
    OSf, {{sigmaf, sigmaF[[b]]}, {tof, toF[[b]]}, {freqf,
      freq[[d]]}, {ampf, ampF[[b]]}, {thetaf, 0}, {OSf, 0}}, t,
    MaxIterations -> 2000];
  sigmaF[[b]] = Fittingnow["BestFitParameters"][[1]][[2]];
  toF[[b]] = Fittingnow["BestFitParameters"][[2]][[2]];
  freqF[[b]] = Fittingnow["BestFitParameters"][[3]][[2]];
  ampF[[b]] = Fittingnow["BestFitParameters"][[4]][[2]];
  \[Theta]F[[b]] = Fittingnow["BestFitParameters"][[5]][[2]];
  OSF[[b]] = Fittingnow["BestFitParameters"][[5]][[2]];
  , {b, 1, numberoffiles}]

  (*First try at means and standard deviations, including non-fitted
  traces*)
  sigmaMean = Mean[Abs[sigmaF]]

```

```

sigmaSD = StandardDeviation[Abs[sigmaF]]
toMean = Mean[toF]
toSD = StandardDeviation[toF]
freqMean = Mean[freqF];
freqSD = StandardDeviation[freqF];
ampMean = Mean[Abs[ampF]]
ampSD = StandardDeviation[Abs[ampF]]
thetaMean = Mean[\[Theta]F];
thetaSD = StandardDeviation[\[Theta]F];
OSMean = Mean[OSF];
OSSD = StandardDeviation[OSF];

(*checking for bad fits*)
sigmaBad = Table[0 i, {i, amin, amax}];
ampBad = Table[0 i, {i, amin, amax}];
ampBVal = Table[0 i, {i, amin, amax}];
Bad = Table[0 i, {i, amin, amax}];
Clear[b]
Do[
  sigmaBad[[b]] =
    If[Abs[Abs[sigmaF[[b]]] - sigmaMean] < 3*sigmaSD, 0, b];
  ampBad[[b]] = If[Abs[Abs[ampF[[b]]] - ampMean] < 3*ampSD, 0, b];
  ampBVal[[b]] =
    If[Abs[Abs[ampF[[b]]] - ampMean] < 3*ampSD, 0, ampF[[b]]];
  Bad[[b]] = Sqrt[sigmaBad[[b]]*ampBad[[b]]
    , {b, 1, numberoffiles}]

(*Plotting out data to check fit*)
b = 2 ;
ModGausFit[t_] =
  ampF[[b]] Exp[-0.5 (t - toF[[b]])^2/(sigmaF[[b]]^2)] \[Epsilon][[
    b]] Cos[2 \[Pi] freqF[[b]] (t - \[Theta]F[[b]])];
Show[ListLinePlot[datashort[[b]], PlotStyle -> Blue,
  PlotRange -> All],
  Plot[ModGausFit[t], {t, 1, enc[[d]] - bnc[[d]]}, PlotStyle -> Green,
  PlotRange -> All]]

(*Removing bad traces*)
Clear[b]
Bad;
(*Bad[[85]]=85
Bad[[86]]=86*)
sigmaF2 = sigmaF;
toF2 = toF;
freqF2 = freqF;
ampF2 = ampF;
\[Theta]F2 = \[Theta]F;
OSF2 = OSF;

```

```

datashort2 = datashort;

Do[
  bi = numberoffiles - b + 1;
  sigmaF2 = If[Bad[[bi]] > 0, Delete[sigmaF2, Bad[[bi]]], sigmaF2];
  toF2 = If[Bad[[bi]] > 0, Delete[toF2, Bad[[bi]]], toF2];
  freqF2 = If[Bad[[bi]] > 0, Delete[freqF2, Bad[[bi]]], freqF2];
  ampF2 = If[Bad[[bi]] > 0, Delete[ampF2, Bad[[bi]]], ampF2];
  \[Theta]F2 =
  If[Bad[[bi]] > 0, Delete[\[Theta]F2, Bad[[bi]]], \[Theta]F2];
  OSF2 = If[Bad[[bi]] > 0, Delete[OSF2, Bad[[bi]]], OSF2];
  datashort2 =
  If[Bad[[bi]] > 0, Delete[datashort2, Bad[[bi]]], datashort2];
  , {b, 1, numberoffiles}]
NoOfTraces = Length[sigmaF2]
ampF2

(*Redefining means etc*)
sigmaMean2 = Mean[Abs[sigmaF2]]
sigmaSD2 = StandardDeviation[Abs[sigmaF2]]
toMean2 = Mean[toF2];
toMean2 + bnc[[d]]
toSD2 = StandardDeviation[toF2]
freqMean2 = Mean[freqF2];
freqSD2 = StandardDeviation[freqF2];
ampMean2 = Mean[Abs[ampF2]]
ampSD2 = StandardDeviation[Abs[ampF2]]
thetaMean2 = Mean[\[Theta]F2];
thetaSD2 = StandardDeviation[\[Theta]F2];
OSMean2 = Mean[OSF2];
OSSD2 = StandardDeviation[OSF2];

(*Digital Demodulation*)
Clear[b]

cmin = 1;
cmax = NoOfTraces;

period = Table[0 i, {i, cmin, cmax}];
DemodLength = Table[0 i, {i, cmin, cmax}];

Do[
  tmax[[b]];
  period[[b]] =
  Round[1/freqF2[[b]]/2];(*manually change for traces that are ng*)

  DemodLength[[b]] = IntegerPart[tmax[[b]]/period[[b]]];
  Demod = Table[0 j + 0 i, {i, 1, numberoffiles}, {j, 1, tmax[[b]]}];

```

```

Demod2 =
  Table[0 j + 0 i, {i, 1, numberoffiles}, {j, 1, DemodLength[[b]]}];
AverageDemod =
  Table[0 j + 0 i, {i, 1, numberoffiles}, {j, 1, DemodLength[[b]]}];
, {b, cmin, cmax}
Do[
  Do[Demod[[b]][[
    t]] = \[Epsilon] [[b]] Cos[
      2 \[Pi] freqF2[[b]] (t - \[Theta]F2[[b]])] datashort2[[b]][[t]];
    , {t, 1, tmax[[b]]}];
  Do[
    Demod2[[b]][[n]] =
      Sum[Demod[[b]][[
        t]], {t, (n - 1)*period[[b]] + 1, (n)*period[[b]]}];
    , {n, 1, DemodLength[[b]] - 1}];
    , {b, cmin, cmax}];

  (*Now Finding Area*)
Clear[b]

Area = Table[0 i, {i, cmin, cmax}];
AreaSq = Table[0 i, {i, cmin, cmax}];

Do[Trapind = Table[0 i, {i, 1, Dimensions[Demod2[[b]]][[1]]}];
  TrapindSq = Table[0 i, {i, 1, Dimensions[Demod2[[b]]][[1]]}];
  Do[
    Trapind[[z]] =
      period[[b]] * .5 (Abs[Demod2[[b]][[z]]] +
        Abs[Demod2[[b]][[z + 1]]]);
    TrapindSq[[z]] =
      period[[b]] * .5 (Demod2[[b]][[z]]^2 + Demod2[[b]][[z + 1]]^2);
    , {z, 1, Dimensions[Demod2[[b]]][[1]] - 1}];
  Area[[b]] =
    Sum[Trapind[[z]], {z, 1, Dimensions[Demod2[[b]]][[1]] - 1}];
  AreaSq[[b]] =
    Sum[TrapindSq[[z]], {z, 1, Dimensions[Demod2[[b]]][[1]] - 1}];
    , {b, cmin, cmax}];
Trapind;
Area^2;
AreaSq
AreaSD = StandardDeviation[Area]
AreaMean = Mean[Area]^2;
AreaSqMean = Mean[AreaSq]
AreaSqSD = StandardDeviation[AreaSq]
Max[AreaSq]

(*For Averaging Everything and putting fit to Final data*)
Clear[b]

```

```

Do[AverageDemod[[1]][[t]] = Abs[Demod2[[1]][[t]]], {t, 1,
  Min[DemodLength]};
Do[
  Do[
    AverageDemod[[b]][[t]] =
      Abs[Demod2[[b]][[t]]] + AverageDemod[[b - 1]][[t]];
    , {t, 1, DemodLength[[b]] - 1};
  , {b, cmin + 1, cmax}];

AverageDemod2 = Abs[AverageDemod]/cmax;

ampavguess = Max[Abs[AverageDemod2[[cmax]]]]
ampavpos = Position[Abs[AverageDemod2[[cmax]]], ampavguess][[1, 1]];

Fittingave =
  NonlinearModelFit[AverageDemod2[[cmax]],
    AmpDemodFit Exp[-0.5 (t -
      todemod)^2/(sigmademod)^2], {{AmpDemodFit,
      ampavguess}, {todemod, ampavpos}, {sigmademod,
      sigmaMean2/Median[period]}}, t, MaxIterations -> 2000];
AmpDF = Fittingave["BestFitParameters"][[1]][[2]]
toDF = Fittingave["BestFitParameters"][[2]][[2]]
sigmaDF = Fittingave["BestFitParameters"][[3]][[2]]

(*Checking Fit*)
DemodFit[t_] = AmpDF Exp[-0.5 (t - toDF)^2/sigmaDF^2];
Show[ListLinePlot[AverageDemod2[[cmax]]],
  Plot[DemodFit[t], {t, 1, 100}, PlotRange -> All, PlotStyle -> Green]]

```

D.3 ac Stark Shift Code

```

(* For investigating agreement between actual energy splitting and
acS6 *)
Off[ClebschGordan::tri ]
Off[ClebschGordan::phy]

(* Constants *)

c = 2.99792458*^8;(* in m/s *)
\[Epsilon]0 =
  1/(c*c*4 \[Pi]*10^-7); (* in F/m *)
kB = 1.3806503*^-23 ;(* J/K *)

h = 6.62606876*^-34;(* in J.]s *)
hbar = h/(2 \[Pi]);
e = 1.602*^-19;(* in C *)

```

```

gf = .5;
mRb = 1.45*10^-25; (*mass of Rb87 in kg*)

(* Mean Frequencies for 87 *)
(*\[Nu][F_, F1_, J_, J1_];*)
\[Nu]mean[
  1/2] = 377.1074635*^12;
kmean[1/2] = 2 \[Pi]/(795*10^-9);
\[Nu]mean[3/2] = 384.2304844685*^12;
(*Frequencies of individual F-F1 transitions in 87*)
(*P1/2 \
transition*)
\[Nu]87[F_, F1_, 1/2, 1/2] =
  377.1074635*^12 -
  3.417*^9/2*(F (F + 1) - 3/2 (3/2 + 1) - 1/2 (1/2 + 1)) +
  406.1*^6/2*(F1 (F1 + 1) - 3/2 (3/2 + 1) - 1/2 (1/2 + 1));
(*P3/2 transition*)
\[Nu]87[F_, F1_, 1/2, 3/2] =
  384.2304844685*^12 -
  3.417*^9/2*(F (F + 1) - 3/2 (3/2 + 1) - 1/2 (1/2 + 1)) +
  84.8*^6/2*(F1 (F1 + 1) - 3/2 (3/2 + 1) - 3/2 (3/2 + 1));
(* Linewidths for 87Rb, assume equal for 85 and 87 Rb: *)

\[CapitalGamma]1[1/2] = 2*\[Pi]*5.746*^6;
\[CapitalGamma]1[3/2] = 2*\[Pi]*6.065*^6;
\[Nu]gs[F_, F1_] =
  -(\[Nu]87[F1, 2, 1/2, 1/2] - \[Nu]87[F, 2, 1/2, 1/2]);

Trec = hbar^2 kmean[1/2]^2/(mRb kB) (*recoil temp*)

(* Reduced matrix coeffs, assume equal for 85 and 87 Rb: *)
redME[1/2, 1/2] = 0.5*2.537*^-29; (* in C*m *)
redME[1/2, 3/2] = 3.584*^-29;
redME[3/2, 1/2] = 0;
redME[3/2, 3/2] =
  0; (*NOT sure about this one!!*)

(* Transition amplitude from
(F,mF,J) to (F1,mF1,J1), polarisation q): *)
ME87[F_, mF_, F1_, mF1_, q_, J_, J1_] =
  Sqrt[(2*1/2 + 1)*(2 F1 + 1)*(2 F + 1)]*
  ThreeJSymbol[{F, mF}, {1,
    q}, {F1, -mF1}]*(SixJSymbol[{J, J1, 1}, {F1, F, 3/2}])*(redME[J,
    J1] + redME[J1, J])*(-1)^(mF1 + 3/2 +
    J); (*Matrix Elements <gf|erq|gi>*)
A[F_, mF_, F1_, mF1_, q_, J_, J1_, F2_, mF2_, qsc_,
  J2_] = (ME87[F2, mF2, F1, mF1, qsc, J2, J1]*
  ME87[F, mF, F2, mF2, q, J, J2]);

```

```

A2[F_, mF_, F1_, mF1_, q_, J_, J1_, F2_, mF2_, qsc_,
  J2_] = (ME87[F1, mF1, F2, mF2, qsc, J1, J2]*
  ME87[F, mF, F2, mF2, q, J, J2]);
A3[F_, mF_, F1_, mF1_, q_, J_, J1_, F2_, mF2_, qsc_,
  J2_] = (Conjugate[ME87[F1, mF1, F2, mF2, qsc, J1, J2]]*
  ME87[F, mF, F2, mF2, q, J, J2]);
(*For scattering count*)

F = 2;
mF = -2;
F2 = 1;
mF2 = -1;
F3 = 1;
q = 1;
J = 1/2;
J1 = 1/2;
Sqrt[Abs[ME87[F, mF, F2, mF2, 0, J, J1]]^2 +
  Abs[ME87[F, mF, F3, mF2, 0, J, J1]]^2];
Abs[ME87[2, -2, 1, -1, 1, J, J1]];
Abs[ME87[2, -2, 2, -1, 1, J, J1]];
Abs[ME87[2, -2, 1, -1, 1, J, J1]]/Abs[ME87[2, -2, 2, -1, 1, J, J1]];
Sqrt[Abs[ME87[2, -2, 1, -1, 1, J, J1]]^2 +
  Abs[ME87[2, -2, 2, -1, 1, J,
    J1]]^2] (*Dipole transition strength for Omega_c assuming right
  circ pol and from F=2 level, mF=-2*)
Sqrt[
  Abs[ME87[2, -2, 2, -2, 0, J,
    J1]]^2] (*Dipole transition strength for Omega_c assuming lin pol
  from F=2 level, mF=-2*)
(*etc*)

(*Laser properties*)
q = 1; (*Polarisation*)

J = 1/2; (*Always 1/2 for S ground state*)
\[Nu]laser[\[Delta]laser_]
= \[Nu]87[2, 1, 1/2,
  1/2] + \[Delta]laser; (* second term is detuning, this determines
  the actual frequency of the laser from a state (1st two terms are the
  F levels and then ??) *)

Intensity = 1; (*2.5*10^2*); (* power/(\[Pi] waist^2)*); (* W/m^2 *)
(* energy shifts in MHz for 87*)
F = 2;
mF = 1;
\[CapitalDelta]E21 [\[Delta]laser_] =
  Intensity/(2 c \[Epsilon]0 h^2 10^6) *
  Sum[(ME87[F, mF, F1, mF1, q, J,

```

```

J1))^2/((\[Nu]laser\[Delta]laser] - \[Nu]87[F, F1, J,
J1])), {J1, 1/2, 3/2}, {F1, 0, 3}, {mF1, -F1, F1}];
(*etc*)

\[Delta]2\[Delta]laser_] =
  Abs\[CapitalDelta]E21\[Delta]laser] -
\[CapitalDelta]E20\[Delta]laser]]; (*energy splitting from mF=0 to
mF=1 levels of F=2 ground state in MHz*)
(*etc*)

(*Laser properties without Rotating Wave Approximation Th2 p .235*)

q = 1; (*Polarisation*)

J = 1/2; (*Always 1/2 for S ground state*)
\[Nu]laser\[Delta]laser_]
= \[Nu]87[2, 1, 1/2, 1/2] + \[Delta]laser;
(* second term is detuning, this determines
the actual frequency of the laser from a state (1st two terms are the
F levels) *)

Intensity =
  1; (*2.5*10^2*); (* power/(\[Pi] waist^2*); (* W/m^2 *) \
(*Shouldn't this be 2*power ??, see Foot p .200*)
(* energy shifts in \
MHz for 87*)
F = 2;
mF = 1;
\[CapitalDelta]E21rwa \[Delta]laser_] =
  Intensity/(2 c \[Epsilon]0 h^2 10^6) *(
  Sum[(ME87[F, mF, F1, mF1, q, J,
J1))^2/((\[Nu]laser\[Delta]laser] - \[Nu]87[F, F1, J,
J1])), {J1, 1/2, 3/2}, {F1, 0, 3}, {mF1, -F1, F1}] +
  Sum[(ME87[F, mF, F1, mF1, q, J,
J1))^2/((\[Nu]laser\[Delta]laser] + \[Nu]87[F, F1, J,
J1])), {J1, 1/2, 3/2}, {F1, 0, 3}, {mF1, -F1, F1}]);
(*etc*)

(*Calculating Scattering Rate for no Rotating Wave Approximation*)
J1 = 1/2;
J = 1/2;
q = 1;
F = 2;
mF = 0;
F = 2;
\[CapitalGamma]sc20rwa\[Delta]laser_] =
  1/(6 \[Pi] \[Epsilon]0^2 hbar^3 c^4) * (Sum[(2 \[Pi] (\[Nu]laser[
\[Delta]laser] - \[Nu]gs[F,

```

```

F1]))^3*(Sum[(A3[F, mF, F1, mF1, q, J, J1, F2, mF2, qsc,
  J2]/(2 \[Pi] (-\[Nu]laser\[Delta]laser] + \[Nu]87[F, F2,
    1/2, J2])) +
A3[F, mF, F1, mF1, q, J, J1, F2, mF2, qsc,
  J2]/(2 \[Pi] (\[Nu]laser\[Delta]laser] + \[Nu]87[F, F2,
    1/2, J2]))), {qsc, -1, 1}, {J2, 1/2, 3/2}, {F2, 0,
  3}, {mF2, -F2, F2}]^2, {F1, 1, 2}, {mF1, -F1, F1}]);
(*etc*)

(*scattering / ac Stark shifts*)
\[CapitalGamma]act11nrwa\[Delta]laser_] = \[CapitalGamma]sc11nrwa[
\[Delta]laser]/\[Delta]rwa\[Delta]laser];
(*etc*)

```

Bibliography

- [1] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414(6862):413–418, 2001. [vii](#), [12](#), [22](#)
- [2] P. W. Shor. Algorithms for Quantum Computation : Discrete Logarithms and Factoring. In *35th Annual Symposium on Foundations of Computer Science, 1994 Proceedings*, pages 124–134, 1994. [vii](#), [12](#)
- [3] V. Giovannetti, S. Lloyd, and L. Maccone. Quantum-enhanced measurements: beating the standard quantum limit. *Science*, 306(5700):1330–1336, November 2004. [vii](#), [16](#)
- [4] 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, January 2011. [vii](#), [29](#), [31](#), [97](#), [102](#), [107](#), [118](#), [180](#), [237](#), [238](#)
- [5] J. J. Longdell, E. Fraval, M. J. Sellars, and N. B. Manson. Stopped Light with Storage Times Greater than One Second Using Electromagnetically Induced Transparency in a Solid. *Physical Review Letters*, 95(6):63601, 2005. [vii](#), [21](#), [31](#), [163](#)
- [6] E. Saglamyurek, N. Sinclair, J. Jin, J. Slater, D. Oblak, F. Bussi eres, M. George, R. Ricken, W. Sohler, and W. Tittel. Conditional Detection of Pure Quantum States of Light after Storage in a Tm-Doped Waveguide. *Physical Review Letters*, 108(8):083602, February 2012. [vii](#), [26](#)
- [7] 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(4):218–221, March 2010. [vii](#), [22](#), [95](#), [96](#), [103](#)
- [8] M. Bonarota, J.-L. Le Gou et, and T. Chaneli ere. Highly multimode storage in a crystal. *New Journal of Physics*, 13(1):013013, January 2011. [vii](#), [26](#), [31](#)
- [9] M. Hosseini, B. M. Sparkes, G. H etet, J. J. Longdell, P. K. Lam, and B. C. Buchler. Coherent optical pulse sequencer for quantum applications. *Nature*, 461(7261):241–245, September 2009. [vii](#), [29](#), [64](#), [66](#), [118](#), [145](#)
- [10] G. Campbell, M. Hosseini, B. M. Sparkes, P. K. Lam, and B. C. Buchler. Time- and frequency-domain polariton interference. *New Journal of Physics*, 14(3):033022, March 2012. [vii](#), [29](#), [63](#), [118](#), [130](#), [133](#), [137](#)
- [11] H. J. Kimble. The quantum internet. *Nature*, 453(7198):1023–30, June 2008. [vii](#), [12](#)
- [12] Y.-F. Chen, C.-Y. Wang, S.-H. Wang, and I. A. Yu. Low-Light-Level Cross-Phase-Modulation Based on Stored Light Pulses. *Physical Review Letters*, 96(4):043603, February 2006. [vii](#), [135](#), [137](#)

-
- [13] Y.-H. Chen, M.-J. Lee, I.-C. Wang, S. Du, Y.-F. Chen, Y.-C. Chen, and I. A. Yu. Coherent Optical Memory with High Storage Efficiency and Large Fractional Delay. *Physical Review Letters*, 110(8):083601, February 2013. [vii](#), [21](#), [31](#), [163](#), [180](#), [198](#)
- [14] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden. Quantum cryptography. *Reviews of Modern Physics*, 74(1):145–195, March 2002. [5](#), [12](#)
- [15] A. Einstein. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik*, 323(6):132–148, 1905. [6](#)
- [16] N. Bohr. On the constitution of atoms and molecules. *Philosophical Magazine Series 6*, 26(153):476–502, September 1913. [6](#)
- [17] H. P. Robertson. The Uncertainty Principle. *Physical Review*, 34(1):163–164, July 1929. [6](#)
- [18] J. G. Rarity, P. C. M. Owens, and P. R. Tapster. Quantum Random-number Generation and Key Sharing. *Journal of Modern Optics*, 41(12):2435–2444, December 1994. [6](#), [15](#)
- [19] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, October 1982. [6](#), [10](#)
- [20] G. I. Taylor. Interference fringes with feeble light. *Proceedings of the Cambridge Philosophical Society*, 15(1):114–115, 1909. [7](#)
- [21] P. Grangier, G. Roger, and A. Aspect. Experimental Evidence for a Photon Anti-correlation Effect on a Beam Splitter: A New Light on Single-Photon Interferences. *Europhysics Letters*, 1(4):173–179, February 1986. [7](#), [15](#)
- [22] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle. Observation of Interference Between Two Bose Condensates. *Science*, 275(5300):637–641, January 1997. [7](#)
- [23] I. Marcikic, H. de Riedmatten, W. Tittel, V. Scarani, H. Zbinden, and N. Gisin. Time-bin entangled qubits for quantum communication created by femtosecond pulses. *Physical Review A*, 66(6):062308, December 2002. [8](#)
- [24] A. P. Lund and T. C. Ralph. Nondeterministic gates for photonic single-rail quantum logic. *Physical Review A*, 66(3):032307, September 2002. [8](#)
- [25] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, January 2001. [8](#), [13](#), [14](#)
- [26] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47(10):777–780, 1935. [8](#), [69](#), [110](#)
- [27] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters. Teleportating an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels. *Physical Review Letters*, 79(13):1895–1899, 1993. [8](#), [12](#)
- [28] A. K. Ekert. Quantum cryptography based on Bells theorem. *Physical Review Letters*, 67(6):661–663, August 1991. [8](#), [10](#)

-
- [29] A. Aspect, P. Grangier, and G. Roger. Experimental Tests of Realistic Local Theories via Bell's Theorem. *Physical Review Letters*, 47(7):460–463, August 1981. 8
- [30] P. G. Kwiat, K. Mattle, H. Weinfurter, A. Zeilinger, A. V. Sergienko, and Y. Shih. New High-Intensity Source of Polarization-Entangled Photon Pairs. *Physical Review Letters*, 75(24):4337–4341, December 1995. 8
- [31] H. J. Kimble and D. F. Walls. Squeezed states of the electromagnetic field: Introduction to feature issue. *Journal of the Optical Society of America B*, 4(10):1449–1450, 1987. 9
- [32] R. E. Slusher, L. W. Hollberg, B. Yurke, J. C. Mertz, and J. F. Valley. Observation of Squeezed States Generated by Four-Wave Mixing in an Optical Cavity. *Physical Review Letters*, 55(22):2409–2412, November 1985. 9
- [33] M. S. Stefszky. *Generation and Detection of Low-Frequency Squeezing for Gravitational-Wave Detection*. PhD thesis, The Australian National University, 2012. 9
- [34] T. Eberle, S. Steinlechner, J. Bauchrowitz, V. Händchen, H. Vahlbruch, M. Mehmet, H. Müller-Ebhardt, and R. Schnabel. Quantum Enhancement of the Zero-Area Sagnac Interferometer Topology for Gravitational Wave Detection. *Physical Review Letters*, 104(25):251102, June 2010. 9
- [35] M. S. Stefszky, C. M. Mow-Lowry, S. S. Y. Chua, D. A. Shaddock, B. C. Buchler, H. Vahlbruch, A. Khalaidovski, R. Schnabel, P. K. Lam, and D. E. McClelland. Balanced homodyne detection of optical quantum states at audio-band frequencies and below. *Classical and Quantum Gravity*, 29(14):145015, July 2012. 9
- [36] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984. 9
- [37] C. Shannon. Communication Theory of Secrecy Systems. *Bell System Technical Journal*, 28(4):656–715, 1949. 10
- [38] P. W. Shor and J. Preskill. Simple proof of security of the BB84 quantum key distribution protocol. *Physical Review Letters*, 85(2):441–444, 2000. 10
- [39] C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin. Experimental Quantum Cryptography. *Journal of Cryptology*, 5(1):3–28, 1992. 10
- [40] R. Ursin, F. Tiefenbacher, T. Schmitt-Manderbach, H. Weier, T. Scheidl, M. Lindenthal, B. Blauensteiner, T. Jennewein, J. Perdigues, P. Trojek, B. Ömer, M. Fürst, M. Meyenburg, J. Rarity, Z. Sodnik, C. Barbieri, H. Weinfurter, and A. Zeilinger. Entanglement-based quantum communication over 144 km. *Nature Physics*, 3(7):481–486, June 2007. 10
- [41] A. Fedrizzi, R. Ursin, T. Herbst, M. Nespoli, R. Prevedel, T. Scheidl, F. Tiefenbacher, T. Jennewein, and A. Zeilinger. High-fidelity transmission of entanglement over a high-loss free-space channel. *Nature Physics*, 5(6):389–392, May 2009. 10

- [42] D. Stucki, N. Walenta, F. Vannel, R. T. Thew, N. Gisin, H. Zbinden, S. Gray, C. R. Towery, and S. Ten. High rate, long-distance quantum key distribution over 250 km of ultra low loss fibres. *New Journal of Physics*, 11(7):075003, July 2009. [10](#)
- [43] S. Wang, W. Chen, J.-F. Guo, Z.-Q. Yin, H.-W. Li, Z. Zhou, G.-C. Guo, and Z.-F. Han. 2-GHz clock quantum key distribution over 260 km of standard telecom fiber. *Optics Letters*, 37(6):1008–1010, March 2012. [10](#)
- [44] S. L. Braunstein and P. van Loock. Quantum information with continuous variables. *Reviews of Modern Physics*, 77(2):513–577, June 2005. [10](#)
- [45] D. Stucki, N. Gisin, O. Guinnard, G. Ribordy, and H. Zbinden. Quantum key distribution over 67 km with a plug&play system. *New Journal of Physics*, 4:41–41, July 2002. [10](#), [19](#)
- [46] B. C. Jacobs and J. D. Franson. Quantum cryptography in free space. *Optics Letters*, 21(22):1854–1856, November 1996. [10](#)
- [47] M. Aspelmeyer, T. Jennewein, M. Pfennigbauer, W. R. Leeb, and A. Zeilinger. Long-Distance Quantum Communication with Entangled Photons Using Satellites. *IEEE Journal of Selected Topics in Quantum Electronics*, 9(6):1541–1551, November 2003. [10](#)
- [48] H.-J. Briegel, W. Dür, J. Cirac, and P. Zoller. Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication. *Physical Review Letters*, 81(26):5932–5935, December 1998. [12](#), [167](#)
- [49] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Nature*, 390(6660):575–579, 1997. [12](#)
- [50] H. Weinfurter. Experimental Bell-State Analysis. *Europhysics Letters*, 25(8):559–564, March 1994. [12](#)
- [51] M. Żukowski, A. Zeilinger, M. Horne, and A. K. Ekert. Event-ready-detectors Bell experiment via entanglement swapping. *Physical Review Letters*, 71(26):4287–4290, December 1993. [12](#)
- [52] J.-W. Pan, Z.-B. Chen, C.-Y. Lu, H. Weinfurter, A. Zeilinger, and M. Żukowski. Multiphoton entanglement and interferometry. *Reviews of Modern Physics*, 84(2):777–838, May 2012. [12](#)
- [53] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Physical Review Letters*, 76(16):3221–3224, 1997. [12](#), [19](#)
- [54] A. M. Lance, T. Symul, W. P. Bowen, B. C. Sanders, and P. K. Lam. Tripartite Quantum State Sharing. *Physical Review Letters*, 92(17):177903, April 2004. [12](#)
- [55] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Review of Modern Physics*, 83(1):33–80, March 2011. [12](#), [16](#), [17](#), [237](#)

-
- [56] R. Prevedel, P. Walther, F. Tiefenbacher, P. Böhi, R. Kaltenbaek, T. Jennewein, and A. Zeilinger. High-speed linear optics quantum computing using active feed-forward. *Nature*, 445(7123):65–69, January 2007. [13](#), [14](#)
- [57] 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*, 28:212–219, 1996. [12](#)
- [58] L. K. Grover. Quantum Mechanics Helps in Searching for a Needle in a Haystack. *Physical Review Letters*, 79(2):325–328, July 1997. [12](#)
- [59] S. Lloyd. Universal Quantum Simulators. *Science*, 273(5278):1073–1078, August 1996. [12](#)
- [60] R. L. Rivest, A. Shamir, and L. M. Adleman. Cryptographic communications system and method. *US Patent*, page US3657476, 1970. [14](#)
- [61] D. P. DiVincenzo. Two-bit gates are universal for quantum computation. *Physical Review A*, 51(2):1015–1022, February 1995. [14](#)
- [62] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Deterministic quantum teleportation of atomic qubits. *Nature*, 429(6993):737–739, June 2004. [14](#)
- [63] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a Fundamental Quantum Logic Gate. *Physical Review Letters*, 75(25):4714–4717, December 1995. [14](#)
- [64] M. Riebe, H. Häffner, C. F. Roos, W. Hänsel, J. Benhelm, G. P. T. Lancaster, T. W. Körber, C. Becher, F. Schmidt-Kaler, D. F. V. James, and R. Blatt. Deterministic quantum teleportation with atoms. *Nature*, 429(6993):734–737, June 2004. [14](#)
- [65] E. Fraval, M. J. Sellars, and J. J. Longdell. Dynamic Decoherence Control of a Solid-State Nuclear-Quadrupole Qubit. *Physical Review Letters*, 95(3):030506, July 2005. [14](#)
- [66] C. K. Hong, Z. Y. Ou, and L. Mandel. Measurement of subpicosecond time intervals between two photons by interference. *Physical Review Letters*, 59(18):2044–2046, November 1987. [14](#)
- [67] D. Gottesman and I. L. Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402(November):390–393, 1999. [14](#)
- [68] N. Lütkenhaus, J. Calsamiglia, and K.-A. Suominen. Bell measurements for teleportation. *Physical Review A*, 59(5):3295–3300, May 1999. [14](#)
- [69] H. J. Briegel and R. Raussendorf. Persistent Entanglement in Arrays of Interacting Particles. *Physical Review Letters*, 86(5):910–913, January 2001. [14](#), [117](#), [133](#)
- [70] R. Raussendorf and H. J. Briegel. A One-Way Quantum Computer. *Physical Review Letters*, 86(22):5188–5191, May 2001. [14](#), [117](#), [133](#)

-
- [71] P. Walther, K. J. Resch, T. Rudolph, E. Schenck, H. Weinfurter, V. Vedral, M. Aspelmeyer, and A. Zeilinger. Experimental one-way quantum computing. *Nature*, 434(7030):169–176, 2005. 14, 117, 133
- [72] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn. Linear optical quantum computing with photonic qubits. *Reviews of Modern Physics*, 79(1):135–174, January 2007. 14
- [73] C. R. Myers and R. Laflamme. Linear Optics Quantum Computation : an Overview. *Arxiv preprint arXiv:quant-ph/0512104v1*, 2005. 14
- [74] A. Kuhn, M. Hennrich, and G. Rempe. Deterministic Single-Photon Source for Distributed Quantum Networking. *Physical Review Letters*, 89(6):067901, July 2002. 15
- [75] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. Continuous generation of single photons with controlled waveform in an ion-trap cavity system. *Nature*, 431(7012):1075–1078, October 2004. 15
- [76] P. Michler, A. Kiraz, C. Becher, W. V. Schoenfeld, P. M. Petroff, Lidong Zhang, E. Hu, and A. Imamoglu. A quantum dot single-photon turnstile device. *Science*, 290(5500):2282–2285, December 2000. 15
- [77] C. Santori, M. Pelton, G. Solomon, Y. Dale, and Y. Yamamoto. Triggered Single Photons from a Quantum Dot. *Physical Review Letters*, 86(8):1502–1505, February 2001. 15
- [78] C. Kurtsiefer, S. Mayer, P. Zarda, and H. Weinfurter. Stable Solid-State Source of Single Photons. *Physical Review Letters*, 85(2):290–293, July 2000. 15
- [79] C. Santori, D. Fattal, J. Vuckovic, G. S. Solomon, and Y. Yamamoto. Single-photon generation with InAs quantum dots. *New Journal of Physics*, 6(1):89, July 2004. 15
- [80] B. Lounis and M. Orrit. Single-photon sources. *Reports on Progress in Physics*, 68(5):1129–1179, May 2005. 15
- [81] C. K. Hong and L. Mandel. Experimental realization of a localized one-photon state. *Physical Review Letters*, 56(1):58–60, January 1986. 15
- [82] J. J. Bollinger, W. M. Itano, D. J. Wineland, and D. J. Heinzen. Optimal frequency measurements with maximally correlated states. *Physical Review A*, 54(6):R4649–R4652, December 1996. 15
- [83] S. F. Huelga, C. Macchiavello, T. Pellizzari, A. K. Ekert, M. B. Plenio, and J. I. Cirac. Improvement of Frequency Standards with Quantum Entanglement. *Physical Review Letters*, 79(20):3865–3868, November 1997. 15
- [84] A. N. Boto, P. Kok, D. S. Abrams, S. L. Braunstein, C. P. Williams, and J. P. Dowling. Quantum Interferometric Optical Lithography: Exploiting Entanglement to Beat the Diffraction Limit. *Physical Review Letters*, 85(13):2733–2736, September 2000. 15
- [85] M. D’Angelo, M. Chekhova, and Y. Shih. Two-Photon Diffraction and Quantum Lithography. *Physical Review Letters*, 87(1):013602, June 2001. 15

-
- [86] R. Jozsa, D. S. Abrams, J. P. Dowling, and C. P. Williams. Quantum clock synchronization based on shared prior entanglement. *Physical Review Letters*, 85(9):2010–2013, August 2000. 15
- [87] V. Giovannetti, S. Lloyd, and L. Maccone. Quantum-enhanced positioning and clock synchronization. *Nature*, 412(6845):417–419, July 2001. 15
- [88] D. V. Strekalov, A. V. Sergienko, D. N. Klyshko, and Y. H. Shih. Observation of Two-Photon Ghost Interference and Diffraction. *Physical Review Letters*, 74(18):3600–3603, May 1995. 15
- [89] T. B. Pittman, Y. H. Shih, D. V. Strekalov, and A. V. Sergienko. Optical imaging by means of two-photon quantum entanglement. *Physical Review A*, 52(5):R3429–R3432, November 1995. 15
- [90] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen. Spin squeezing and reduced quantum noise in spectroscopy. *Physical Review A*, 46(11):R6797–R6800, December 1992. 16
- [91] W. Wasilewski, K. Jensen, H. Krauter, J. J. Renema, M. V. Balabas, and E. S. Polzik. Quantum Noise Limited and Entanglement-Assisted Magnetometry. *Physical Review Letters*, 104(13):133601, March 2010. 16
- [92] M. I. Kolobov. The spatial behavior of nonclassical light. *Reviews of Modern Physics*, 71(5):1539–1589, October 1999. 16
- [93] L. A. Lugiato, A. Gatti, and E. Brambilla. Quantum imaging. *Journal of Optics B: Quantum and Semiclassical Optics*, 4(3):S176–S183, June 2002. 16
- [94] M. Varnava, D. E. Browne, and T. Rudolph. Loss Tolerance in One-Way Quantum Computation via Counterfactual Error Correction. *Physical Review Letters*, 97(12):120501, September 2006. 16
- [95] E. Wigner. On the Quantum Correction For Thermodynamic Equilibrium. *Physical Review*, 40(5):749–759, June 1932. 16
- [96] F. Grosshans and P. Grangier. Quantum cloning and teleportation criteria for continuous quantum variables. *Physical Review A*, 64(1):010301, June 2001. 16, 57, 112
- [97] N. J. Cerf, O. Krüger, P. Navez, R. F. Werner, and M. M. Wolf. Non-Gaussian Cloning of Quantum Coherent States is Optimal. *Physical Review Letters*, 95(7):070501, August 2005. 16, 56, 112
- [98] P. Grangier, J. A. Levenson, and J.-P. Poizat. Quantum non-demolition measurements in optics. *Nature*, 396(6711):537–542, 1998. 16, 56
- [99] T. C. Ralph and P. K. Lam. Teleportation with Bright Squeezed Light. *Physical Review Letters*, 81(25):5668–5671, December 1998. 16, 56, 57
- [100] N. Gisin and R. Thew. Quantum communication. *Nature Photonics*, 1(3):165–171, March 2007. 17, 237

- [101] M. Fiorentino, S. M. Spillane, R. G. Beausoleil, T. D. Roberts, P. Battle, and M. W. Munro. Spontaneous parametric down-conversion in periodically poled KTP waveguides and bulk crystals. *Optics Express*, 15(12):7479–7488, June 2007. [17](#)
- [102] F. Wolfgramm, X. Xing, A. Cerè, A. Predojević, A. M. Steinberg, and M. W. Mitchel. Bright filter-free source of indistinguishable photon pairs. *Optics Express*, 16(22):18145–18151, 2008. [17](#)
- [103] H. Zhang, X.-M. Jin, J. Yang, H.-N. Dai, S.-J. Yang, T.-M. Zhao, J. Rui, Y. He, X. Jiang, F. Yang, G.-S. Pan, Z.-S. Yuan, Y. Deng, Z.-B. Chen, X.-H. Bao, S. Chen, B. Zhao, and J.-W. Pan. Preparation and storage of frequency-uncorrelated entangled photons from cavity-enhanced spontaneous parametric downconversion. *Nature Photonics*, 5(10):628–632, September 2011. [17](#), [164](#)
- [104] D. V. Vasilyev, I. V. Sokolov, and E. S. Polzik. Quantum memory for images: A quantum hologram. *Physical Review A*, 77(2):020302(R), February 2008. [18](#), [118](#)
- [105] 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(19):190503, May 2007. [18](#), [118](#), [139](#), [237](#)
- [106] O. A. Collins, S. D. Jenkins, A. Kuzmich, and T. A. B. Kennedy. Multiplexed Memory-Insensitive Quantum Repeaters. *Physical Review Letters*, 98(6):060502, February 2007. [18](#), [118](#), [139](#), [237](#)
- [107] X. Li, P. L. Voss, J. Chen, J. E. Sharping, and P. Kumar. Storage and long-distance distribution of telecommunications-band polarization entanglement generated in an optical fiber. *Optics Letters*, 30(10):1201–1203, May 2005. [19](#)
- [108] R. M. Gingrich, P. Kok, H. Lee, F. Vatan, and J. P. Dowling. All Linear Optical Quantum Memory Based on Quantum Error Correction. *Physical Review Letters*, 91(21):217901, November 2003. [19](#)
- [109] T. B. Pittman and J. D. Franson. Cyclical quantum memory for photonic qubits. *Physical Review A*, 66(6):062302, December 2002. [19](#)
- [110] T. Tanabe, M. Notomi, E. Kuramochi, A. Shinya, and H. Taniyama. Trapping and delaying photons for one nanosecond in an ultrasmall high-Q photonic-crystal nanocavity. *Nature Photonics*, 1(1):49–52, January 2007. [19](#)
- [111] S. Kuhr, S. Gleyzes, C. Guerlin, J. Bernu, U. B. Hoff, S. Deleglise, S. Osnaghi, M. Brune, J.-M. Raimond, S. Haroche, E. Jacques, P. Bosland, and B. Visentin. Ultrahigh finesse Fabry-Perot superconducting resonator. *Applied Physics Letters*, 90(16):164101, 2007. [19](#)
- [112] A. I. Lvovsky, B. C. Sanders, and W. Tittel. Optical quantum memory. *Nature Photonics*, 3(12):706–714, December 2009. [19](#)
- [113] A. S. Parkins, P. Marte, P. Zoller, and H. J. Kimble. Synthesis of arbitrary quantum states via adiabatic transfer of Zeeman coherence. *Physical Review Letters*, 71(19):3095–3098, November 1993. [19](#)

-
- [114] J. Ye, D. W. Vernooy, and H. J. Kimble. Trapping of Single Atoms in Cavity QED. *Physical Review Letters*, 83(24):4987–4990, December 1999. [19](#)
- [115] M. Hijlkema, B. Weber, H. P. Specht, S. C. Webster, A. Kuhn, and G. Rempe. A single-photon server with just one atom. *Nature Physics*, 3(4):253–255, March 2007. [19](#)
- [116] K. M. Fortier, S. Y. Kim, M. J. Gibbons, P. Ahmadi, and M. S. Chapman. Deterministic Loading of Individual Atoms to a High-Finesse Optical Cavity. *Physical Review Letters*, 98(23):233601, June 2007. [19](#)
- [117] H. P. Specht, C. Nölleke, A. Reiserer, M. Uphoff, E. Figueroa, S. Ritter, and G. Rempe. A single-atom quantum memory. *Nature*, 473(7346):190–193, May 2011. [19](#)
- [118] L. V. Hau, S. E. Harris, Z. Dutton, and C. H. Behroozi. Light speed reduction to 17 metres per second in an ultracold atomic gas. *Nature*, 397(6720):594–598, 1999. [20](#), [163](#)
- [119] R. Zhang, S. R. Garner, and L. V. Hau. Creation of Long-Term Coherent Optical Memory via Controlled Nonlinear Interactions in Bose-Einstein Condensates. *Physical Review Letters*, 103(23):233602, December 2009. [20](#), [163](#), [173](#)
- [120] M. Fleischhauer and M. D. Lukin. Dark-State Polaritons in Electromagnetically Induced Transparency. *Physical Review Letters*, 84(22):5094–5097, May 2000. [20](#), [60](#)
- [121] C. Liu, Z. Dutton, C. H. Behroozi, and L. V. Hau. Observation of coherent optical information storage in an atomic medium using halted light pulses. *Nature*, 409(6819):490–493, 2001. [20](#), [96](#), [163](#)
- [122] D. F. Phillips, A. Fleischhauer, A. Mair, R. L. Walsworth, and M. Lukin. Storage of Light in Atomic Vapor. *Physical Review Letters*, 86(5):783–786, January 2001. [20](#), [96](#)
- [123] A. V. Turukhin, V. S. Sudarshanam, M. S. Shahriar, J. A. Musser, B. S. Ham, and P. R. Hemmer. Observation of Ultraslow and Stored Light Pulses in a Solid. *Physical Review Letters*, 88(2):023602, December 2001. [20](#)
- [124] N. B. Phillips, A. V. Gorshkov, and I. Novikova. Optimal light storage in atomic vapor. *Physical Review A*, 78(2):023801, 2008. [21](#), [96](#), [103](#), [151](#)
- [125] 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(12):123601, March 2007. [21](#)
- [126] 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(7069):837–841, December 2005. [21](#), [96](#), [99](#)
- [127] 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(7069):833–836, December 2005. [21](#)

- [128] D. Akamatsu, K. Akiba, and M. Kozuma. Electromagnetically Induced Transparency with Squeezed Vacuum. *Physical Review Letters*, 92(20):203602, May 2004. [21](#)
- [129] J. Appel, E. Figueroa, D. Korystov, M. Lobino, and A. I. Lvovsky. Quantum Memory for Squeezed Light. *Physical Review Letters*, 100(9):093602, March 2008. [21](#), [96](#)
- [130] K. Honda, D. Akamatsu, M. Arikawa, Y. Yokoi, K. Akiba, S. Nagatsuka, T. Tanimura, A. Furusawa, and M. Kozuma. Storage and Retrieval of a Squeezed Vacuum. *Physical Review Letters*, 100(9):093601, March 2008. [21](#)
- [131] M. Arikawa, K. Honda, D. Akamatsu, S. Nagatsuka, K. Akiba, A. Furusawa, and M. Kozuma. Quantum memory of a squeezed vacuum for arbitrary frequency sidebands. *Physical Review A*, 81(2):021605, February 2010. [21](#), [164](#), [194](#)
- [132] K. S. Choi, H. Deng, J. Laurat, and H. J. Kimble. Mapping photonic entanglement into and out of a quantum memory. *Nature*, 452(7183):67–71, March 2008. [21](#)
- [133] E. E. Mikhailov, T. Horrom, N. Belcher, and I. Novikova. Performance of a prototype atomic clock based on linlin coherent population trapping resonances in Rb atomic vapor. *Journal of the Optical Society of America B*, 27(3):417–422, February 2010. [21](#)
- [134] J. Vanier. Atomic clocks based on coherent population trapping: a review. *Applied Physics B*, 81(4):421–442, July 2005. [21](#), [58](#)
- [135] Marlan O. Scully and Michael Fleischhauer. High-sensitivity magnetometer based on index-enhanced media. *Physical Review Letters*, 69(9):1360–1363, August 1992. [21](#)
- [136] G. Hétet, A. Peng, M. T. Johnsson, M. T. L. Hsu, O. Glöckl, P. K. Lam, H. A. Bachor, and J. J. Hope. Erratum: Squeezing and entanglement delay using slow light [Phys. Rev. A 71, 033809 (2005)]. *Physical Review A*, 74(5):059902(E), November 2006. [21](#)
- [137] G. Hétet, A. Peng, M. T. Johnsson, J. J. Hope, and P. K. Lam. Characterization of electromagnetically-induced-transparency-based continuous-variable quantum memories. *Physical Review A*, 77(1):12323, January 2008. [21](#), [49](#)
- [138] M. T. L. Hsu, G. Hétet, O. Glöckl, J. J. Longdell, B. C. Buchler, H.-A. Bachor, and P. K. Lam. Quantum Study of Information Delay in Electromagnetically Induced Transparency. *Physical Review Letters*, 97(18):183601, November 2006. [21](#)
- [139] E. Figueroa, F. Vewinger, J. Appel, and A. I. Lvovsky. Decoherence of electromagnetically induced transparency in atomic vapor. *Optics Letters*, 31(17):2625–2627, September 2006. [21](#), [49](#), [98](#)
- [140] J. Nunn, K. F. Reim, K. C. Lee, V. O. Lorenz, B. J. Sussman, I. A. Walmsley, and D. Jaksch. Multimode Memories in Atomic Ensembles. *Physical Review Letters*, 101(26):260502, December 2008. [21](#), [28](#), [63](#)

-
- [141] 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(1):011401(R), January 2007. [21](#)
- [142] 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(5):053603, July 2011. [22](#), [31](#), [96](#), [102](#), [104](#)
- [143] K. Reim, J. Nunn, X.-M. Jin, P. Michelberger, T. Champion, D. England, K. Lee, W. Kolthammer, N. Langford, and I. Walmsley. Multipulse Addressing of a Raman Quantum Memory: Configurable Beam Splitting and Efficient Readout. *Physical Review Letters*, 108(26):263602, June 2012. [22](#)
- [144] D. G. England, P. S. Michelberger, T. F. M. Champion, K. F. Reim, K. C. Lee, M. R. Sprague, X.-M. Jin, N. K. Langford, W. S. Kolthammer, J. Nunn, and I. A. Walmsley. High-fidelity polarization storage in a gigahertz bandwidth quantum memory. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(12):124008, June 2012. [22](#)
- [145] C. H. van der Wal, M. D. Eisaman, A. André, R. L. Walsworth, D. F. Phillips, A. S. Zibrov, and M. D. Lukin. Atomic memory for correlated photon states. *Science*, 301(5630):196–200, July 2003. [22](#), [96](#), [99](#)
- [146] A. Kuzmich, W. P. Bowen, A. D. Boozer, A. Boca, C. W. Chou, L.-M. Duan, and H. J. Kimble. Generation of nonclassical photon pairs for scalable quantum communication with atomic ensembles. *Nature*, 423(6941):731–734, June 2003. [23](#)
- [147] J. Laurat, H. de Riedmatten, D. Felinto, C.-W. Chou, E. W. Schomburg, and H. J. Kimble. Efficient retrieval of a single excitation stored in an atomic ensemble. *Optics Express*, 14(15):6912–6918, July 2006. [23](#)
- [148] J. Simon, H. Tanji, J. K. Thompson, and V. Vuletić. Interfacing Collective Atomic Excitations and Single Photons. *Physical Review Letters*, 98(18):183601, May 2007. [23](#)
- [149] X.-H. Bao, A. Reingruber, P. Dietrich, J. Rui, A. Dück, T. Strassel, L. Li, N.-L. Liu, B. Zhao, and J.-W. Pan. Efficient and long-lived quantum memory with cold atoms inside a ring cavity. *Nature Physics*, 8(7):517–521, May 2012. [23](#), [163](#), [164](#)
- [150] V. Balić, D. A. Braje, P. Kolchin, G. Y. Yin, and S. E. Harris. Generation of Paired Photons with Controllable Waveforms. *Physical Review Letters*, 94(18):184601, May 2005. [23](#), [232](#)
- [151] C. W. Chou, H. de Riedmatten, D. Felinto, S. V. Polyakov, S. J. van Enk, and H. J. Kimble. Measurement-induced entanglement for excitation stored in remote atomic ensembles. *Nature*, 438(7069):828–832, December 2005. [23](#)
- [152] K. S. Choi, A. Goban, S. B. Papp, S. J. van Enk, and H. J. Kimble. Entanglement of spin waves among four quantum memories. *Nature*, 468(7322):412–416, November 2010. [23](#)
- [153] E. L. Hahn. Spin Echoes. *Physical Review*, 80(4):580–594, November 1950. [23](#)

- [154] N. A. Kurnit, I. D. Abella, and S. R. Hartmann. Observation of a Photon Echo. *Physical Review Letters*, 13(19):567–568, November 1964. [23](#)
- [155] J. Ruggiero, J.-L. Le Gouët, C. Simon, and T. Chanelière. Why the two-pulse photon echo is not a good quantum memory protocol. *Physical Review A*, 79(5):053851, May 2009. [24](#)
- [156] N. Sangouard, C. Simon, M. Afzelius, and N. Gisin. Analysis of a quantum memory for photons based on controlled reversible inhomogeneous broadening. *Physical Review A*, 75(3):032327, March 2007. [24](#), [28](#)
- [157] P. M. Ledingham, W. R. Naylor, J. J. Longdell, S. E. Beavan, and M. J. Sellars. Nonclassical photon streams using rephased amplified spontaneous emission. *Physical Review A*, 81(1):012301, January 2010. [24](#)
- [158] S. E. Beavan, M. P. Hedges, and M. J. Sellars. Demonstration of Photon-Echo Rephasing of Spontaneous Emission. *Physical Review Letters*, 109(9):093603, August 2012. [24](#)
- [159] P. M. Ledingham, W. R. Naylor, and J. J. Longdell. Experimental Realization of Light with Time-Separated Correlations by Rephasing Amplified Spontaneous Emission. *Physical Review Letters*, 109(9):093602, August 2012. [24](#)
- [160] T. W. Mossberg. Time-domain frequency-selective optical data storage. *Optics Letters*, 7(2):77–79, 1982. [25](#)
- [161] M. K. Kim and R. Kachru. Multiple-bit long-term data storage by backward-stimulated echo in $\text{Eu}^{3+}:\text{YAlO}_3$. *Optics Letters*, 14(9):423–425, 1989. [25](#)
- [162] H. Lin, T. Wang, and T. W. Mossberg. Demonstration of 8-Gbit/in.² areal storage density based on swept-carrier frequency-selective optical memory. *Optics Letters*, 20(15):1658–1660, 1995. [25](#)
- [163] K. D. Merkel, R. Krishna Mohan, Z. Cole, T. Chang, A. Olson, and W. R. Babbitt. Multi-Gigahertz radar range processing of baseband and RF carrier modulated signals in $\text{Tm}:\text{YAG}$. *Journal of Luminescence*, 107(1-4):62–74, May 2004. [25](#)
- [164] B. S. Ham. Atom phase-locked coherence conversion using optical locking for ultralong photon storage beyond the spin T_2 constraint. *New Journal of Physics*, 14(1):013003, January 2012. [25](#)
- [165] D. L. McAuslan, P. M. Ledingham, W. R. Naylor, S. E. Beavan, M. P. Hedges, M. J. Sellars, and J. J. Longdell. Photon-echo quantum memories in inhomogeneously broadened two-level atoms. *Physical Review A*, 84(2):022309, August 2011. [25](#)
- [166] V. Damon, M. Bonarota, A. Louchet-Chauvet, T. Chanelière, and J.-L. Le Gouët. Revival of silenced echo and quantum memory for light. *New Journal of Physics*, 13(9):093031, September 2011. [25](#)
- [167] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin. Multimode quantum memory based on atomic frequency combs. *Physical Review A*, 79(5):052329, May 2009. [25](#), [26](#)

-
- [168] H. de Riedmatten, M. Afzelius, M. U. Staudt, C. Simon, and N. Gisin. A solid-state light-matter interface at the single-photon level. *Nature*, 456(7223):773–777, December 2008. [26](#)
- [169] A. Amari, A. Walther, M. Sabooni, M. Huang, S. Kröll, M. Afzelius, I. Usmani, B. Lauritzen, N. Sangouard, H. de Riedmatten, and N. Gisin. Towards an efficient atomic frequency comb quantum memory. *Journal of Luminescence*, 130(9):1579–1585, September 2010. [26](#)
- [170] E. Saglamyurek, N. Sinclair, J. Jin, J. A. Slater, D. Oblak, F. Bussi eres, M. George, R. Ricken, W. Sohler, and W. Tittel. Broadband waveguide quantum memory for entangled photons. *Nature*, 469(7331):512–515, January 2011. [26](#), [31](#)
- [171] C. Clausen, I. Usmani, F. Bussi eres, N. Sangouard, M. Afzelius, H. de Riedmatten, and N. Gisin. Quantum storage of photonic entanglement in a crystal. *Nature*, 469(7331):508–511, January 2011. [26](#)
- [172] I. Usmani, C. Clausen, F. Bussi eres, N. Sangouard, M. Afzelius, and N. Gisin. Heralded quantum entanglement between two crystals. *Nature Photonics*, 6(4):234–237, March 2012. [26](#)
- [173] M. Afzelius, I. Usmani, A. Amari, B. Lauritzen, A. Walther, C. Simon, N. Sangouard, J. Min ar, H. de Riedmatten, N. Gisin, and S. Kröll. Demonstration of Atomic Frequency Comb Memory for Light with Spin-Wave Storage. *Physical Review Letters*, 104(4):040503, January 2010. [26](#)
- [174] M. G undoĝan, M. Mazzera, P. M. Ledingham, M. Cristiani, and H. de Riedmatten. Coherent storage of temporally multimode light using a spin-wave atomic frequency comb memory. *New Journal of Physics*, 15(4):045012, April 2013. [26](#)
- [175] 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(17):173601, October 2001. [26](#)
- [176] M. Nilsson and S. Kröll. Solid state quantum memory using complete absorption and re-emission of photons by tailored and externally controlled inhomogeneous absorption profiles. *Optics Communications*, 247(4-6):393–403, March 2005. [28](#)
- [177] B. Kraus, W. Tittel, N. Gisin, M. Nilsson, S. Kröll, and J. I. Cirac. Quantum memory for nonstationary light fields based on controlled reversible inhomogeneous broadening. *Physical Review A*, 73(2):020302(R), February 2006. [28](#), [167](#)
- [178] A. L. Alexander, J. J. Longdell, M. J. Sellars, and N. B. Manson. Coherent information storage with photon echoes produced by switching electric fields. *Journal of Luminescence*, 127(1):94–97, 2007. [28](#)
- [179] B. Lauritzen, J. Min ar, 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(8):080502, February 2010. [28](#)
- [180] G. H etet, 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(2):23601, 2008. [28](#)

-
- [181] M. P. Hedges, J. J. Longdell, Y. Li, and M. J. Sellars. Efficient quantum memory for light. *Nature*, 465(7301):1052–1056, June 2010. 28, 112
- [182] B. M. Sparkes. *Development of a Three-Level Gradient Echo Memory*. Honours thesis, The Australian National University, 2008. 29, 97, 167
- [183] 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. 29, 97
- [184] D. A. Steck. Rubidium 87 D line data. *available online at <http://steck.us/alkalidata> (revision 2.0.1, 2 May 2008)*. 29, 101, 154, 156, 241
- [185] B. C. Buchler, M. Hosseini, G. Hétet, B. M. Sparkes, and P. K. Lam. Precision spectral manipulation of optical pulses using a coherent photon echo memory. *Optics Letters*, 35(7):1091–3, April 2010. 29, 118, 128, 138
- [186] J. J. Longdell, G. Hétet, P. K. Lam, and M. J. Sellars. Analytic treatment of controlled reversible inhomogeneous broadening quantum memories for light using two-level atoms. *Physical Review A*, 78(3):32337, September 2008. 29, 62, 63
- [187] F. Carreño and M. A. Antón. Coherent control of light pulses stored in a Gradient Echo Memory. *Optics Communications*, 284(12):3154–3159, June 2011. 29, 118
- [188] F. Carreño and M. A. Antón. Gradient echo memory in a tripod-like dense atomic medium. *Optics Communications*, 283(23):4787–4795, December 2010. 29, 118
- [189] V. Boyer, A. M. Marino, R. C. Pooser, and P. D. Lett. Entangled images from four-wave mixing. *Science*, 321(5888):544–547, July 2008. 30, 97
- [190] Q. Glorieux, J. B. Clark, N. V. Corzo, and P. D. Lett. Generation of pulsed bipartite entanglement using four-wave mixing. *New Journal of Physics*, 14(12):123024, December 2012. 30, 97, 232
- [191] I. H. Agha, G. Messin, and P. Grangier. Generation of pulsed and continuous-wave squeezed light with ^{87}Rb vapor. *Optics Express*, 18(5):4198–205, March 2010. 30, 97
- [192] I. H. Agha, C. Giarmatzi, Q. Glorieux, T. Coudreau, P. Grangier, and G. Messin. Time-resolved detection of relative-intensity squeezed nanosecond pulses in an ^{87}Rb vapor. *New Journal of Physics*, 13(4):043030, April 2011. 30, 97
- [193] V. Boyer, C. F. McCormick, E. Arimondo, and P. D. Lett. Ultraslow Propagation of Matched Pulses by Four-Wave Mixing in an Atomic Vapor. *Physical Review Letters*, 99(14):143601, October 2007. 30, 97
- [194] A. M. Marino, R. C. Pooser, V. Boyer, and P. D. Lett. Tunable delay of Einstein-Podolsky-Rosen entanglement. *Nature*, 457(7231):859–862, 2009. 30, 97, 139
- [195] R. M. Camacho, P. K. Vudyaetu, and J. C. Howell. Four-wave-mixing stopped light in hot atomic rubidium vapour. *Nature Photonics*, 3(2):103–106, January 2009. 30, 97, 103

-
- [196] B. Julsgaard, J. Sherson, J. I. Cirac, J. Flurášek, and E. S. Polzik. Experimental demonstration of quantum memory for light. *Nature*, 432(7016):482–486, 2004. [30](#), [97](#)
- [197] K. Jensen, W. Wasilewski, H. Krauter, T. Fernholz, B. M. Nielsen, M. Owari, M. B. Plenio, A. Serafini, M. M. Wolf, and E. S. Polzik. Quantum memory for entangled continuous-variable states. *Nature Physics*, 7(1):13–16, November 2010. [30](#), [97](#)
- [198] J. F. Sherson, H. Krauter, R. K. Olsson, B. Julsgaard, K. Hammerer, I. Cirac, and E. S. Polzik. Quantum teleportation between light and matter. *Nature*, 443(7111):557–560, October 2006. [30](#)
- [199] Z. Zhu, D. J. Gauthier, and R. W. Boyd. Stored light in an optical fiber via stimulated Brillouin scattering. *Science*, 318(5857):1748–1750, December 2007. [31](#)
- [200] P. C. Maurer, G. Kucsko, C. Latta, L. Jiang, N. Y. Yao, S. D. Bennett, F. Pastawski, D. Hunger, N. Chisholm, M. Markham, D. J. Twitchen, J. I. Cirac, and M. D. Lukin. Room-Temperature Quantum Bit Memory Exceeding One Second. *Science*, 336(6086):1283–1286, June 2012. [31](#)
- [201] M. Steger, K. Saeedi, M. L. W. Thewalt, J. J. L. Morton, H. Riemann, N. V. Abrosimov, P. Becker, and H.-J. Pohl. Quantum Information Storage for over 180 s Using Donor Spins in a ^{28}Si ”Semiconductor Vacuum”. *Science*, 336(6086):1280–1283, June 2012. [31](#)
- [202] C. Clausen, F. Bussi eres, M. Afzelius, and N. Gisin. Quantum Storage of Heralded Polarization Qubits in Birefringent and Anisotropically Absorbing Materials. *Physical Review Letters*, 108(19):190503, May 2012. [31](#)
- [203] L. Allen and J. H. Eberly. *Optical Resonance and Two-Level Atoms*. Dover Publications, 1987. [35](#), [51](#)
- [204] C. J. Foot. *Atomic Physics*. Oxford University Press, 2005. [35](#), [45](#), [50](#), [98](#), [152](#), [154](#), [156](#), [158](#), [185](#), [241](#)
- [205] M. Fox. *Quantum Optics: An Introduction*. Oxford University Press, 2006. [35](#)
- [206] H.-A. Bachor and T. C. Ralph. *A Guide to Experiments in Quantum Optics*. Wiley, second edition, 2004. [35](#), [41](#), [271](#)
- [207] R. Loudon. *The Quantum Theory of Light*. Oxford Science Publications, third edition, 2000. [35](#), [101](#), [160](#)
- [208] R. R. Puri. *Mathematical Methods of Quantum Optics*. Springer, 2001. [35](#), [42](#)
- [209] P. Meystre and M. Sargent III. *Elements of Quantum Optics*. Springer, fourth edition, 2007. [35](#)
- [210] G. H etet. *Quantum Memories for Continuous Variable States of Light in Atomic Ensembles*. PhD thesis, The Australian National University, 2008. [35](#)
- [211] B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics*. John Wiley & Sons., Inc., 1991. [43](#), [185](#)

-
- [212] A. Peng, M. Johnsson, and J. J. Hope. Pulse retrieval and soliton formation in a nonstandard scheme for dynamic electromagnetically induced transparency. *Physical Review A*, 71(3):33817, March 2005. 45
- [213] B. R. Mollow. Power Spectrum of Light Scattered by Two-Level Systems. *Physical Review*, 188(5):1969–1975, December 1969. 48
- [214] 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(3):033804, September 2007. 49, 65
- [215] A. A. Clerk, M. H. Devoret, S. M. Girvin, Florian Marquardt, and R. J. Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Reviews of Modern Physics*, 82(2):1155–1208, April 2010. 56
- [216] Q. Y. He, M. D. Reid, E. Giacobino, J. Cviklinski, and P. D. Drummond. Dynamical oscillator-cavity model for quantum memories. *Physical Review A*, 79(2):022310, February 2009. 56
- [217] S. L. Braunstein, C. A. Fuchs, and H. J. Kimble. Criteria for continuous-variable quantum teleportation. *Journal of Modern Optics*, 47(2-3):267–278, February 2000. 56
- [218] K. Hammerer, M. M. Wolf, E. S. Polzik, and J. I. Cirac. Quantum Benchmark for Storage and Transmission of Coherent States. *Physical Review Letters*, 94(15):150503, April 2005. 56
- [219] J.-Ph. Poizat, J.-F. Roch, and P. Grangier. Characterization of quantum non-demolition measurements in optics. *Annales de Physique*, 19(3):265–297, 1994. 56
- [220] U. Fano. Effects of Configuration Interaction on Intensities and Phase Shifts. *Physical Review*, 124(6):1866–1878, December 1961. 58, 208
- [221] S. Autler and C. Townes. Stark Effect in Rapidly Varying Fields. *Physical Review*, 100(2):703–722, October 1955. 58, 208
- [222] C. Cohen-Tannoudji and S. Reynaud. Simultaneous saturation of two atomic transitions sharing a common level. *Journal of Physics B: Atomic and Molecular Physics*, 10(12):2311–2331, August 1977. 58, 208
- [223] M. Fleischhauer, A. Imamoglu, and J. P. Marangos. Electromagnetically induced transparency: Optics in coherent media. *Reviews of Modern Physics*, 77(2):633–673, 2005. 58
- [224] S. A. Moiseev and N. M. Arslanov. Efficiency and fidelity of photon-echo quantum memory in an atomic system with longitudinal inhomogeneous broadening. *Physical Review A*, 78(2):23803, August 2008. 63, 128
- [225] 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(20):203601, November 2008. 63, 64, 65
- [226] T. Ahola, J. Hu, and E. Ikonen. A digital control system for the iodine stabilized HeNe laser. *Review of Scientific Instruments*, 69(5):1934–1937, 1998. 69

-
- [227] W. Z. Zhao, J. E. Simsarian, L. A. Orozco, and G. D. Sprouse. A computer-based digital feedback control of frequency drift of multiple lasers. *Review of Scientific Instruments*, 69(11):3737–3740, 1998. 69
- [228] J. A. Smith, X. Chu, W. Huang, J. Wiig, and A. T. Brown. LabVIEW-based laser frequency stabilization system with phase-sensitive detection servo loop for Doppler LIDAR applications. *Optical Engineering*, 47(11):114201, 2008. 69
- [229] T. Eom, H. Choi, and S. Lee. Frequency stabilization of an internal mirror He-Ne laser by digital control. *Review of Scientific Instruments*, 73(1):221–224, 2002. 69
- [230] D. V. Pantelić, B. M. Panić, and A. G. Kovačević. Digital control of an iodine stabilized He-Ne laser by using a personal computer and a simple electronic system. *Review of Scientific Instruments*, 74(6):3155–3159, 2003. 69
- [231] E. Canuto and A. Rolino. Multi-input digital frequency stabilization of monolithic lasers. *Automatica*, 40(12):2137–2145, December 2004. 69
- [232] K. Matsubara, S. Uetake, H. Ito, Y. Li, K. Hayasaka, and M. Hosokawa. Precise Frequency-Drift Measurement of Extended-Cavity Diode Laser Stabilized with Scanning Transfer Cavity. *Japanese Journal of Applied Physics*, 44(1A):229–230, January 2005. 69
- [233] N. Seymour-Smith, P. Blythe, M. Keller, and W. Lange. Fast scanning cavity offset lock for laser frequency drift stabilization. *Review of Scientific Instruments*, 81(7):075109, July 2010. 69
- [234] A. Rossi, V. Biancalana, B. Mai, and L. Tomassetti. Long-term drift laser frequency stabilization using purely optical reference. *Review of Scientific Instruments*, 73(7):2544–2548, 2002. 69
- [235] G. J. Koch. Automatic laser frequency locking to gas absorption lines. *Optical Engineering*, 42(6):1690–1693, 2003. 69
- [236] S. K. Lee, B. W. Han, and D. Cho. Automatic system to relock a laser frequency to a FabryPerot cavity. *Review of Scientific Instruments*, 76(2):026101, 2005. 69
- [237] L. Dong, W. Yin, W. Ma, and S. Jia. A novel control system for automatically locking a diode laser frequency to a selected gas absorption line. *Measurement Science and Technology*, 18(5):1447–1452, May 2007. 69
- [238] D. Stoler. Equivalence classes of minimum uncertainty packets. *Physical Review D*, 1(12):3217–3219, 1970. 69, 110
- [239] E. Y. C. Lu. Quantum correlations in two-photon amplification. *Letters Al Nuovo Cimento*, 3(14):585–589, 1972. 69, 110
- [240] D. F. Walls and P. Zoller. Reduced Quantum Fluctuations in Resonance Fluorescence. *Physical Review Letters*, 47(10):709–711, 1981. 70
- [241] H. Yuen and J. Shapiro. Optical communication with two-photon coherent states—Part III: quantum measurements realizable with photoemissive detectors. *IEEE Transactions on Information Theory*, 26(1):78–92, 1980. 70

-
- [242] W. P. Bowen, N. Treps, B. C. Buchler, R. Schnabel, T. C. Ralph, H.-A. Bachor, T. Symul, and P. K. Lam. Experimental investigation of continuous-variable quantum teleportation. *Physical Review A*, 67(3):032302, March 2003. 70, 85
- [243] M. Yukawa, R. Ukai, P. van Loock, and A. Furusawa. Experimental generation of four-mode continuous-variable cluster states. *Physical Review A*, 78(1):012301, July 2008. 70
- [244] R. H. Bishop and R. C. Dorf. *Modern Control Systems*. Addison-Wesley, eighth edition, 1998. 71, 80, 252
- [245] H. Nyquist. Regeneration Theory. *Bell System Technical Journal*, 11(3):126–147, 1932. 74, 252
- [246] R. V. Pound. Electronic Frequency Stabilization of Microwave Oscillators. *Review of Scientific Instruments*, 17(11):490–505, 1946. 75, 246, 273
- [247] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser Phase and Frequency Stabilization Using an Optical Resonator. *Applied Physics B: Lasers and Optics*, 31(2):97–105, 1983. 75, 246, 273
- [248] E. B. Hogenauer. An Economical class of digital filters for decimation and interpolation. *IEEE Transactions on Acoustic, Speech, and Signal Processing*, 29(2):155–162, 1981. 77, 248
- [249] A. R. Hambley. *Electrical Engineering. Principles and Applications*. Pearson Prentice Hall, third edition, 2005. 87
- [250] A. Ourjoumtsev, R. Tualle-Brouiri, J. Laurat, and P. Grangier. Generating optical Schrödinger kittens for quantum information processing. *Science*, 312(5770):83–86, April 2006. 90
- [251] H. M. Chrzanowski, J. Bernu, B. M. Sparkes, B. Hage, A.P. Lund, T. C. Ralph, P. K. Lam, and T. Symul. Photon-number discrimination without a photon counter and its application to reconstructing non-Gaussian states. *Physical Review A*, 84(5):050302(R), November 2011. 90
- [252] M. Hosseini, B. M. Sparkes, G. T. Campbell, P. K. Lam, and B. C. Buchler. Storage and manipulation of light using a Raman gradient-echo process. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(12):124004, June 2012. 95, 102, 103, 125, 138, 151, 218
- [253] T. Baluktsian, C. Urban, T. Bublath, H. Giessen, R. Löw, and T. Pfau. Fabrication method for microscopic vapor cells for alkali atoms. *Optics Letters*, 35(12):1950–1952, June 2010. 95
- [254] B. Wu, J. F. Hulbert, E. J. Lunt, K. Hurd, A. R. Hawkins, and H. Schmidt. Slow light on a chip via atomic quantum state control. *Nature Photonics*, 4(11):776–779, September 2010. 95
- [255] M. R. Sprague, D. G. England, A. Abdolvand, J. Nunn, X.-M. Jin, W. S. Kolthammer, M. Barbieri, B. Rigal, P. S. Michelberger, T. F. M. Champion, P. St.J. Russell, and I. A. Walmsley. Efficient optical pumping and high optical depth in a hollow-core

-
- photonic-crystal fibre for a broadband quantum memory. *New Journal of Physics*, 15(5):055013, May 2013. 95
- [256] I. Novikova, A. V. Gorshkov, D. F. Phillips, A. S. Sørensen, M. D. Lukin, and R. L. Walsworth. Optimal Control of Light Pulse Storage and Retrieval. *Physical Review Letters*, 98(24):243602, June 2007. 96, 179
- [257] I. Novikova, N. B. Phillips, and A. V. Gorshkov. Optimal light storage with full pulse-shape control. *Physical Review A*, 78(2):021802(R), August 2008. 96, 118
- [258] M. Hosseini. *Quantum Optical Storage and Processing Using Raman Gradient Echo Memory*. PhD thesis, The Australian National University, 2012. 97, 113, 137
- [259] A. Javan, O. Kocharovskaya, H. Lee, and M. O. Scully. Narrowing of electromagnetically induced transparency resonance in a Doppler-broadened medium. *Physical Review A*, 66(1):013805, July 2002. 98
- [260] M. Erhard, S. Nußmann, and H. Helm. Power broadening and Doppler effects of coherent dark resonances in Rb. *Physical Review A*, 62(6):061802(R), November 2000. 98
- [261] O. Firstenberg, M. Shuker, R. Pugatch, D. Fredkin, N. Davidson, and A. Ron. Theory of thermal motion in electromagnetically induced transparency: Effects of diffusion, Doppler broadening, and Dicke and Ramsey narrowing. *Physical Review A*, 77(4):043830, April 2008. 99
- [262] R. Pugatch, M. Shuker, O. Firstenberg, A. Ron, and N. Davidson. Topological Stability of Stored Optical Vortices. *Physical Review Letters*, 98(20):203601, May 2007. 99, 139
- [263] M. Shuker, O. Firstenberg, R. Pugatch, A. Ron, and N. Davidson. Storing Images in Warm Atomic Vapor. *Physical Review Letters*, 100(22):223601, June 2008. 99, 139
- [264] D. Higginbottom, B. M. Sparkes, M. Rancic, O. Pinel, M. Hosseini, P. K. Lam, and B. C. Buchler. Spatial-mode storage in a gradient-echo memory. *Physical Review A*, 86(2):023801, August 2012. 99
- [265] I. Novikova, Y. Xiao, D. F. Phillips, and R. L. Walsworth. EIT and diffusion of atomic coherence. *Journal of Modern Optics*, 52(16):2381–2390, November 2005. 99
- [266] Y. Xiao, I. Novikova, D. F. Phillips, and R. L. Walsworth. Diffusion-Induced Ramsey Narrowing. *Physical Review Letters*, 96(4):043601, February 2006. 99
- [267] W. Happer. Optical Pumping. *Reviews of Modern Physics*, 44(2):169–249, April 1972. 100
- [268] M. Erhard and H. Helm. Buffer-gas effects on dark resonances: Theory and experiment. *Physical Review A*, 63(4):043813, March 2001. 100, 101
- [269] S. Brandt, A. Nagel, R. Wynands, and D. Meschede. Buffer-gas-induced linewidth reduction of coherent dark resonances to below 50 Hz. *Physical Review A*, 56(2):R1063–R1066, August 1997. 101

- [270] M. D. Rotondaro and G. P. Perram. Collisional broadening and shift of the rubidium D1 and D2 lines ($^5S_{1/2}$ - $^5P_{3/2}$) by rare gases, H₂, D₂, N₂, CH₄ and CF₄. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 57(4):497–507, April 1997. 101
- [271] E. Arimondo. Relaxation processes in coherent-population trapping. *Physical Review A*, 54(3):2216–2223, September 1996. 101
- [272] M. A. Bouchiat and J. Brossel. Relaxation of Optically Pumped Rb Atoms on Paraffin-Coated Walls. *Physical Review*, 147(1):41–54, July 1966. 101
- [273] S. Jiang, X.-M. Luo, L.-Q. Chen, B. Ning, S. Chen, J.-Y. Wang, Z.-P. Zhong, and J.-W. Pan. Observation of prolonged coherence time of the collective spin wave of an atomic ensemble in a paraffin-coated ^{87}Rb vapor cell. *Physical Review A*, 80(6):062303, December 2009. 101
- [274] S. J. Seltzer, P. J. Meares, and M. V. Romalis. Synchronous optical pumping of quantum revival beats for atomic magnetometry. *Physical Review A*, 75(5):051407, May 2007. 101
- [275] M. V. Balabas, K. Jensen, W. Wasilewski, H. Krauter, L. S. Madsen, J. H. Müller, T. Fernholz, and E. S. Polzik. High quality anti-relaxation coating material for alkali atom vapor cells. *Optics Express*, 18(6):5825–5830, March 2010. 101, 151
- [276] M. V. Balabas, T. Karaulanov, M. P. Ledbetter, and D. Budker. Polarized Alkali-Metal Vapor with Minute-Long Transverse Spin-Relaxation Time. *Physical Review Letters*, 105(7):070801, August 2010. 101, 151
- [277] M. L. Citron, H. R. Gray, C. W. Gabel, and C.R. Stroud Jr. Experimental study of power broadening in a two-level atom. *Physical Review A*, 16(4):1507–1512, October 1977. 101
- [278] N. B. Phillips, A. V. Gorshkov, and I. Novikova. Slow light propagation and amplification via electromagnetically induced transparency and four-wave mixing in an optically dense atomic vapor. *Journal of Modern Optics*, 56(18-19):1916–1925, October 2009. 102, 103
- [279] M. D. Lukin, A. B. Matsko, M. Fleischhauer, and M. O. Scully. Quantum Noise and Correlations in Resonantly Enhanced Wave Mixing Based on Atomic Coherence. *Physical Review Letters*, 82(9):1847–1850, March 1999. 102
- [280] D. Higginbottom. *Spatial Multimode Storage in a Gradient Echo Memory*. Honours thesis, The Australian National University, 2011. 102, 143
- [281] C. W. Thiel. *Four-Wave Mixing and its Applications*, 2008. 102
- [282] M. D. Lukin, M. Fleischhauer, A. S. Zibrov, H. G. Robinson, V. L. Velichansky, L. Hollberg, and M. O. Scully. Spectroscopy in Dense Coherent Media: Line Narrowing and Interference Effects. *Physical Review Letters*, 79(16):2959–2962, October 1997. 102
- [283] M. D. Lukin, P. R. Hemmer, M. Löffler, and M. O. Scully. Resonant Enhancement of Parametric Processes via Radiative Interference and Induced Coherence. *Physical Review Letters*, 81(13):2675–2678, September 1998. 102

-
- [284] E. E. Mikhailov, Y. V. Rostovtsev, and G. R. Welch. Group velocity study in hot ^{87}Rb vapour with buffer gas. *Journal of Modern Optics*, 50(15-17):2645–2654, January 2003. 103
- [285] A. B. Matsko, D. V. Strekalov, and L. Maleki. On the dynamic range of optical delay lines based on coherent atomic media. *Optics Express*, 13(6):2210–2223, 2005. 103
- [286] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler. Unconditional room-temperature quantum memory. *Nature Physics*, 7(10):794–798, June 2011. 111, 112, 113
- [287] I. L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *Journal of Modern Optics*, 44(11-12):2455–2467, November 1997. 110
- [288] J. L. O’Brien, G. J. Pryde, A. Gilchrist, D. F. V. James, N. K. Langford, T. C. Ralph, and A. G. White. Quantum Process Tomography of a Controlled-NOT Gate. *Physical Review Letters*, 93(8):080502, August 2004. 110
- [289] M. Lobino, D. Korystov, C. Kupchak, E. Figueroa, B. C. Sanders, and A. I. Lvovsky. Complete characterization of quantum-optical processes. *Science*, 322(5901):563–566, October 2008. 110
- [290] M. Lobino, C. Kupchak, E. Figueroa, and A. I. Lvovsky. Memory for Light as a Quantum Process. *Physical Review Letters*, 102(20):203601, May 2009. 110
- [291] S. Rahimi-Keshari, A. Scherer, A. Mann, A. T. Rezakhani, A. I. Lvovsky, and B. C. Sanders. Quantum process tomography with coherent states. *New Journal of Physics*, 13(1):013006, January 2011. 110
- [292] J. Řeháček, Z. Hradil, and M. Ježek. Iterative algorithm for reconstruction of entangled states. *Physical Review A*, 63(4):040303, March 2001. 113
- [293] S. A. Moiseev and W. Tittel. Temporal compression of quantum-information-carrying photons using a photon-echo quantum memory approach. *Physical Review A*, 82(1):012309, July 2010. 117, 135, 167
- [294] M. Brandt-Pearce, I. Jacobs, J.-H. Lee, and J. K. Shaw. Optimal input Gaussian pulse width for transmission in dispersive nonlinear fibers. *Journal of the Optical Society of America B*, 16(8):1189–1196, 1999. 117
- [295] G. J. Milburn. Quantum Optical Fredkin Gate. *Physical Review Letters*, 62(18):2124–2127, 1989. 117, 133
- [296] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble. Measurement of Conditional Phase Shifts Quantum Logic. *Physical Review Letters*, 75(25):4710–4713, 1995. 117, 133
- [297] E. B. Treacy. Optical Pulse Compression With Diffraction Gratings. *IEEE Journal of Quantum Electronics*, 5(9):454–458, 1969. 117

-
- [298] H. P. Specht, J. Bochmann, M. Mücke, B. Weber, E. Figueroa, D. L. Moehring, and G. Rempe. Phase shaping of single-photon wave packets. *Nature Photonics*, 3(8):469–472, July 2009. 117
- [299] P. Kolchin, C. Belthangady, S. Du, G. Y. Yin, and S. E. Harris. Electro-Optic Modulation of Single Photons. *Physical Review Letters*, 101(10):103601, September 2008. 117
- [300] B. H. Kolner. Active pulse compression using an integrated electro-optic phase modulator. *Applied Physics Letters*, 52(14):1122–1124, 1988. 117
- [301] M. A. Dugan, J. X. Tull, and W. S. Warren. High-resolution acousto-optic shaping of unamplified and amplified femtosecond laser pulses. *Journal of the Optical Society of America B*, 14(9):2348–2358, 1997. 117
- [302] A. M. Weiner. Femtosecond pulse shaping using spatial light modulators. *Review of Scientific Instruments*, 71(5):1929–1960, 2000. 117
- [303] S. Viciani, A. Zavatta, and M. Bellini. Nonlocal modulations on the temporal and spectral profiles of an entangled photon pair. *Physical Review A*, 69(5):053801, May 2004. 117
- [304] A. Pe’er, B. Dayan, A. A. Friesem, and Y. Silberberg. Temporal Shaping of Entangled Photons. *Physical Review Letters*, 94(7):073601, February 2005. 117
- [305] B. Dayan, A. Pe’er, A. A. Friesem, and Y. Silberberg. Two Photon Absorption and Coherent Control with Broadband Down-Converted Light. *Physical Review Letters*, 93(2):023005, July 2004. 117
- [306] B. H. Kolner and M. Nazarathy. Temporal imaging with a time lens. *Optics Letters*, 14(12):630–632, 1989. 117
- [307] M. A. Foster, R. Salem, Y. Okawachi, A. C. Turner-Foster, M. Lipson, and A. L. Gaeta. Ultrafast waveform compression using a time-domain telescope. *Nature Photonics*, 3(10):581–585, September 2009. 117
- [308] C. V. Bennett and B. H. Kolner. Upconversion time microscope demonstrating 103x magnification of femtosecond waveforms. *Optics Letters*, 24(11):783–785, 1999. 117
- [309] C. V. Bennett, R. P. Scott, and B. H. Kolner. Temporal magnification and reversal of 100 Gb/s optical data with an up conversion time microscope. *Applied Physics B: Lasers and Optics*, 65(20):2513–2515, 1994. 117
- [310] J. Azaña. Time-to-frequency conversion using a single time lens. *Optics Communications*, 217(1-6):205–209, March 2003. 117
- [311] M. T. Kauffman, W. C. Banyai, A. A. Godil, and D. M. Bloom. Time-to-frequency converter for measuring picosecond optical pulses. *Applied Physics Letters*, 64(3):270–272, 1994. 117
- [312] H. J. McGuinness, M.G. Raymer, C. J. McKinstrie, and S. Radic. Quantum Frequency Translation of Single-Photon States in a Photonic Crystal Fiber. *Physical Review Letters*, 105(9):093604, August 2010. 117

-
- [313] M. A. Foster, A. C. Turner, R. Salem, M. Lipson, and A. L. Gaeta. Broad-band continuous-wave parametric wavelength conversion in silicon nanowaveguides. *Optics Express*, 15(20):12949–12958, 2007. 117
- [314] D. Kielpinski, J. F. Corney, and H. M. Wiseman. Quantum Optical Waveform Conversion. *Physical Review Letters*, 106(13):130501, March 2011. 117
- [315] 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(6):065029, June 2011. 117
- [316] A. Eckstein, B. Brecht, and C. Silberhorn. A quantum pulse gate based on spectrally engineered sum frequency generation. *Optics Express*, 19(15):13770–13778, July 2011. 117
- [317] M. T. Rakher, L. Ma, M. Davanço, O. Slattery, X. Tang, and K. Srinivasan. Simultaneous Wavelength Translation and Amplitude Modulation of Single Photons from a Quantum Dot. *Physical Review Letters*, 107(8):083602, August 2011. 117
- [318] S. Ramelow, A. Fedrizzi, A. Poppe, N. K. Langford, and A. Zeilinger. Polarization-entanglement-conserving frequency conversion of photons. *Physical Review A*, 85(1):013845, January 2012. 117
- [319] J. Huang and P. Kumar. Observation of Quantum Frequency Conversion. *Physical Review Letters*, 68(14):2153–2156, April 1992. 117
- [320] H. Takesue. Single-photon frequency down-conversion experiment. *Physical Review A*, 82(1):013833, July 2010. 117
- [321] Y. Ding and Z. Y. Ou. Frequency downconversion for a quantum network. *Optics Letters*, 35(15):2591–2593, 2010. 117
- [322] A. P. VanDevender and P. G. Kwiat. High-speed transparent switch via frequency upconversion. *Optics Express*, 15(8):4677–4683, 2007. 117
- [323] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters. Mixed-state entanglement and quantum error correction. *Physical Review A*, 54(5):3824–3851, November 1996. 118
- [324] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera. Quantum Privacy Amplification and the Security of Quantum Cryptography over Noisy Channels. *Physical Review Letters*, 77(13):2818–2821, September 1996. 118
- [325] 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(23):233602, November 2004. 118
- [326] G. Heinze, N. Rentzsch, and T. Halfmann. Multiplexed image storage by electromagnetically induced transparency in a solid. *Physical Review A*, 86(5):053837, November 2012. 118, 140
- [327] D.-S. Ding, J. H. Wu, Z. Y. Zhou, Y. Liu, B.-S. Shi, X.-B. Zou, and G.-C. Guo. Multi-mode quantum image memory based on a cold atomic ensemble. *Arxiv preprint arXiv:1204.1130v3*, 2012. 118, 139

-
- [328] H.-N. Dai, H. Zhang, S.-J. Yang, T.-M. Zhao, J. Rui, Y.-J. Deng, L. Li, N.-L. Liu, S. Chen, X.-H. Bao, X.-M. Jin, B. Zhao, and J.-W. Pan. Holographic Storage of Biphoton Entanglement. *Physical Review Letters*, 108(21):210501, May 2012. 118
- [329] J. Wen, Y.-H. Zhai, S. Du, and M. Xiao. Engineering biphoton wave packets with an electromagnetically induced grating. *Physical Review A*, 82(4):043814, October 2010. 118
- [330] A. W. Brown and M. Xiao. All-optical switching and routing based on an electromagnetically induced absorption grating. *Optics Letters*, 30(7):699–701, 2005. 118
- [331] S A Moiseev and B S Ham. Photon-echo quantum memory with efficient multipulse readings. *Physical Review A*, 70(6):63809, 2004. 118
- [332] M. U. Staudt, S. R. Hastings-Simon, M. Nilsson, M. Afzelius, V. Scarani, R. Ricken, H. Suche, W. Sohler, W. Tittel, and N. Gisin. Fidelity of an Optical Memory Based on Stimulated Photon Echoes. *Physical Review Letters*, 98(11):113601, March 2007. 118
- [333] S.-Y. Lan, A. G. Radnaev, O. A. Collins, D. N. Matsukevich, T. A. B. Kennedy, and A. Kuzmich. A Multiplexed Quantum Memory. *Optics Express*, 17(16):13639–13645, 2009. 118
- [334] G. R. Dennis, J. J. Hope, and M. T. Johnsson. XMDS2: Fast, scalable simulation of coupled stochastic partial differential equations. *Computer Physics Communications*, 184(1):201–208, January 2013. 125, 143
- [335] H. Schmidt and A. Imamoglu. Giant Kerr nonlinearities obtained by electromagnetically induced transparency. *Optics Letters*, 21(23):1936–1938, 1996. 133
- [336] B.-W. Shiao, M.-C. Wu, C.-C. Lin, and Y.-C. Chen. Low-Light-Level Cross-Phase Modulation with Double Slow Light Pulses. *Physical Review Letters*, 106(19):193006, May 2011. 133, 137
- [337] H.-Y. Lo, Y.-C. Chen, P.-C. Su, H.-C. Chen, J.-X. Chen, Y.-C. Chen, I. A. Yu, and Y.-F. Chen. Electromagnetically-induced-transparency-based cross-phase-modulation at attojoule levels. *Physical Review A*, 83(4):041804, April 2011. 135, 137
- [338] L. Zhao, G. Yang, and W. Duan. Manipulating stored images with phase imprinting at low light levels. *Optics Letters*, 37(14):2853, July 2012. 135, 139, 147
- [339] J. Gea-Banacloche. Impossibility of large phase shifts via the giant Kerr effect with single-photon wave packets. *Physical Review A*, 81(4):043823, April 2010. 135
- [340] M. Hosseini, S. Rebić, B. M. Sparkes, J. Twamley, B. C. Buchler, and P. K. Lam. Memory-enhanced noiseless cross-phase modulation. *Light: Science & Applications*, 1(12):e40, December 2012. 136
- [341] S. E. Harris and Y. Yamamoto. Photon Switching by Quantum Interference. *Physical Review Letters*, 81(17):3611–3614, 1998. 137

-
- [342] I. V. Sokolov, M. I. Kolobov, A. Gatti, and L. A. Lugiato. Quantum holographic teleportation. *Optics Communications*, 193(1):175–180, 2001. [139](#)
- [343] A. Gatti, I. V. Sokolov, M. I. Kolobov, and L. A. Lugiato. Quantum fluctuations in holographic teleportation of optical images. *The European Physical Journal D*, 30(1):123–135, July 2004. [139](#)
- [344] T. Yu. Golubeva, YU. M. Golubev, I. V. Sokolov, and M. I. Kolobov. Quantum parallel dense coding of optical images. *Journal of Modern Optics*, 53(5-6):699–711, March 2006. [139](#)
- [345] K. Tordrup, A. Negretti, and K. Mølmer. Holographic Quantum Computing. *Physical Review Letters*, 101(4):040501, July 2008. [139](#)
- [346] M. Malik, O. S. Magana-Loaiza, and R. W. Boyd. Quantum-secured imaging. *Applied Physics Letters*, 101(24):241103, 2012. [139](#)
- [347] P. K. Vudyasetu, R. M. Camacho, and J. C. Howell. Storage and Retrieval of Multimode Transverse Images in Hot Atomic Rubidium Vapor. *Physical Review Letters*, 100(12):123903, March 2008. [139](#), [147](#)
- [348] A. E. Siegman. Hermite-Gaussian functions of complex argument as optical-beam eigenfunctions. *Journal of the Optical Society of America*, 63(9):1093–1094, 1973. [139](#), [147](#)
- [349] D. Yankelev, O. Firstenberg, and P. London. Shape-preserving diffusion of a high-order mode. *Optics Letters*, 38(8):1203–1205, April 2013. [139](#)
- [350] D.-S. Ding, J.-H. Wu, Z.-Y. Zhou, B.-S. Shi, X.-B. Zou, and G.-C. Guo. Frequency-multiplexed image storage and conversion in a cold atomic ensemble. *arXiv preprint arXiv:1205.2741v4*, May 2012. [139](#)
- [351] J. Wu, Y. Liu, D.-S. Ding, Z.-Y. Zhou, B.-S. Shi, and G.-C. Guo. Light storage based on four-wave mixing and electromagnetically induced transparency in cold atoms. *Physical Review A*, 87(1):013845, January 2013. [139](#), [147](#)
- [352] G. Heinze, A. Rudolf, F. Beil, and T. Halfmann. Storage of images in atomic coherences in a rare-earth-ion-doped solid. *Physical Review A*, 81(1):011401(R), January 2010. [139](#)
- [353] Q. Glorieux, J. B. Clark, A. M. Marino, Z. Zhou, and P. D. Lett. Temporally multiplexed storage of images in a Gradient Echo Memory. *Optics Express*, 20(11):12350–12358, 2012. [140](#)
- [354] S. Chu, L. Hollberg, J. E. Bjorkholm, A. Cable, and A. Ashkin. Three-Dimensional Viscous Confinement and Cooling of Atoms by Resonance Radiation Pressure. *Physical Review Letters*, 55(1):48–51, 1985. [154](#)
- [355] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59(23):2631–2634, December 1987. [155](#)
- [356] H. J. Metcalf and P. van der Straten. Laser cooling and trapping of atoms. *Journal of the Optical Society of America B*, 20(5):887–908, 2003. [155](#), [165](#), [215](#)

-
- [357] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *Journal of the Optical Society of America B*, 6(11):2023–2045, November 1989. 155, 156, 187
- [358] S. L. Lukman, M. T. DePue, N. E. Bramall, and D. S. Weiss. Laser cooling at high density in deep far-detuned optical lattices. *Physical Review A*, 59(1):R19–R22, January 1999. 156
- [359] W. Ketterle and N. J. van Druten. Evaporative cooling of trapped atoms. *Advances in Atomic, Molecular, and Optical Physics*, 37:181–236, 1996. 156
- [360] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. Cohen-Tannoudji. Laser Cooling below the One-Photon Recoil Energy by Velocity-Selective Coherent Population Trapping. *Physical Review Letters*, 61(7):826–829, August 1988. 156
- [361] M. Kasevich and S. Chu. Laser Cooling below a Photon Recoil with Three-Level Atoms. *Physical Review Letters*, 69(12):1741–1744, September 1992. 156
- [362] V. Bose. Plancks Gesetz und Lichtquantenhypothese. *Zeitschrift für Physik*, 26(1):178–181, December 1924. 156
- [363] A. Einstein. Quantum theory of monatomic ideal gases. *Sitzungsberichte der Preussischen Akademie der Wissenschaften Physikalisch-Mathematische Klasse*, 3:261, 1924. 156
- [364] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269(5221):198–201, July 1995. 156
- [365] A. Ashkin. Trapping of Atoms by Resonance Radiation Pressure. *Physical Review Letters*, 40(12):729–732, 1978. 156
- [366] J. P. Barrat and C. Cohen-Tannoudji. Étude du pompage optique dans le formalisme de la matrice densité. *Le Journal de Physique et le Radium*, 22(6):329–336, 1961. 158, 169, 238
- [367] J. P. Barrat and C. Cohen-Tannoudji. Élargissement et déplacement des raies de resonance magnétique causes par une excitation optique. *Le Journal de Physique et le Radium*, 22(7):443–450, 1961. 158, 169, 238
- [368] S. Chu, J. E. Bjorkholm, A. Ashkin, and A. Cable. Experimental Observation of Optically Trapped Atoms. *Physical Review Letters*, 57(3):314–318, 1986. 158, 159
- [369] N. Davidson, H. J. Lee, C. S. Adams, M. Kasevich, and S. Chu. Long Atomic Coherent Times in an Optical Dipole Trap. *Physical Review Letters*, 74(8):1311–1314, 1995. 158
- [370] H. J. Lee, C. S. Adams, M. Kasevich, and S. Chu. Raman cooling of atoms in an optical dipole trap. *Physical Review Letters*, 76(15):2658–2661, May 1996. 158, 165, 215
- [371] C. Salomon, J. Dalibard, A. Aspect, H. Metcalf, and C. Cohen-Tannoudji. Channeling atoms in a laser standing wave. *Physical Review Letters*, 59(15):1659–1662, October 1987. 158

-
- [372] C. S. Adams, H. J. Lee, N. Davidson, M. Kasevich, and S. Chu. Evaporative cooling in a crossed dipole trap. *Physical review letters*, 74(18):3577–3580, May 1995. 158
- [373] D. Frese, B. Ueberholz, S. Kuhr, W. Alt, D. Schrader, V. Gomer, and D. Meschede. Single Atoms in an optical dipole trap: towards a deterministic source of cold atoms. *Physical Review Letters*, 85(18):3777–3780, 2000. 158
- [374] K. L. Corwin. *A Circularly-Polarized Optical Dipole Trap and Other Developments in Laser Trapping of Atoms*. PhD thesis, University of Colorado, 1999. 158
- [375] N. P. Robins. Personal communication, 2013. 158
- [376] J. D. Miller, R. A. Cline, and D. J. Heinzen. Far-off-resonance optical trapping of atoms. *Physical Review A*, 47(6):R4567–4570, 1993. 158
- [377] R. A. Cline, J. D. Miller, M. R. Mathews, and D. J. Heinzen. Spin Relaxation of optically trapped atoms by light scattering. *Optics Letters*, 19(3):207–209, 1994. 158, 162
- [378] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. Optical dipole traps for neutral atoms. *Advances in Atomic, Molecular, and Optical Physics*, 42:95–170, 2000. 159, 162, 170, 178
- [379] C. D. Wallace, T. P. Dinneen, K.-Y. N. Tan, T. T. Grove, and P. L. Gould. Isotopic Difference in Trap Loss Collisions of Laser Cooled Rubidium Atoms. *Physical Review Letters*, 69(6):897–900, 1992. 159
- [380] D. Sesko, T. Walker, C. Monroe, A. Gallagher, and C. Wieman. Collisional Losses from a Light-Force Atom Trap. *Physical Review Letters*, 63(9):961–964, 1989. 159
- [381] S. J. M. Kuppens, K. L. Corwin, K. W. Miller, T. E. Chupp, and C. E. Wieman. Loading an optical dipole trap. *Physical Review A*, 62(1):13406, 2000. 159
- [382] D. Döring. *Ein Experiment zum zustandsabhängigen Transport einzelner Atome*. Diplomarbeit in physik, Universität Bonn, 2007. 160
- [383] T. A. Savard, K. M. OHara, and J. E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. *Physical Review A*, 56(2):R1095–R1098, August 1997. 162
- [384] M. E. Gehm, K. M. OHara, T. A. Savard, and J. E. Thomas. Dynamics of noise-induced heating in atom traps. *Physical Review A*, 58(5):3914–3921, November 1998. 162
- [385] B. Zhao, Y.-A. Chen, X.-H. Bao, T. Strassel, C.-S. Chuu, X.-M. Jin, J. Schmiedmayer, Z.-S. Yuan, S. Chen, and J.-W. Pan. A millisecond quantum memory for scalable quantum networks. *Nature Physics*, 5(2):95–99, 2008. 163
- [386] R. Zhao, Y. O. Dudin, S. D. Jenkins, C. J. Campbell, D. N. Matsukevich, T. A. B. Kennedy, and A. Kuzmich. Long-lived quantum memory. *Nature Physics*, 5(2):100–104, 2008. 163
- [387] Y. O. Dudin, S. D. Jenkins, R. Zhao, D. N. Matsukevich, A. Kuzmich, and T. A. B. Kennedy. Entanglement of a photon and an optical lattice spin wave. *Physical Review Letters*, 103(2):020505, 2009. 163

- [388] Y. O. Dudin, R. Zhao, T. A. B. Kennedy, and A. Kuzmich. Light storage in a magnetically dressed optical lattice. *Physical Review A*, 81(4):041805(R), April 2010. [163](#)
- [389] U. Schnorrberger, J. D. Thompson, S. Trotzky, R. Pugatch, N. Davidson, S. Kuhr, and I. Bloch. Electromagnetically induced transparency and light storage in an atomic mott insulator. *Physical Review Letters*, 103(3):033003, 2009. [163](#), [173](#)
- [390] S. Zhang, S. Zhou, M. M. T. Loy, G. K. L. Wong, and S. Du. Optical storage with electromagnetically induced transparency in a dense cold atomic ensemble. *Optics Letters*, 36(23):4530–4532, 2011. [163](#)
- [391] S. Riedl, M. Lettner, C. Vo, S. Baur, G. Rempe, and S. Dürr. Bose-Einstein condensate as a quantum memory for a photonic polarization qubit. *Physical Review A*, 85(2):022318, February 2012. [163](#)
- [392] Y. Sagi, I. Almog, and N. Davidson. Process Tomography of Dynamical Decoupling in a Dense Cold Atomic Ensemble. *Physical Review Letters*, 105(5):053201, July 2010. [164](#)
- [393] F. Yang, T. Mandel, C. Lutz, Z.-S. Yuan, and J.-W. Pan. Transverse mode revival of a light-compensated quantum memory. *Physical Review A*, 83(6):063420, June 2011. [164](#)
- [394] A. G. Radnaev, Y. O. Dudin, R. Zhao, H. H. Jen, S. D. Jenkins, A. Kuzmich, and T. A. B. Kennedy. A quantum memory with telecom-wavelength conversion. *Nature Physics*, 6(11):894–899, September 2010. [164](#)
- [395] A. Kaplan, M. F. Andersen, T. Grünzweig, and N. Davidson. Hyperfine spectroscopy of optically trapped atoms. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(8):R103–R125, August 2005. [164](#)
- [396] C. J. Myatt, N. R. Newbury, R. W. Ghrist, S. Loutzenhiser, and C. E. Wieman. Multiply loaded magneto-optical trap. *Optics Letters*, 21(4):290–292, March 1996. [165](#), [231](#)
- [397] Y.-W. Lin, H.-C. Chou, P. P. Dwivedi, Y.-C. Chen, and I. A. Yu. Using a pair of rectangular coils in the MOT for the production of cold atom clouds with large optical density. *Optics Express*, 16(6):3753–3761, March 2008. [165](#), [172](#)
- [398] T. Peters, Y.-H. Chen, J.-S. Wang, Y.-W. Lin, and I. A. Yu. Optimizing the retrieval efficiency of stored light pulses. *Optics Express*, 17(8):6665–6675, 2009. [165](#), [172](#)
- [399] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell. Behavior of atoms in a compressed magneto-optical trap. *Journal of the Optical Society of America B*, 11(8):1332, August 1994. [165](#), [215](#)
- [400] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin, and D. E. Pritchard. High Densities of Cold Atoms in a Dark Spontaneous-Force Optical Trap. *Physical Review Letters*, 70(15):2253–2256, 1993. [165](#)
- [401] S. Zhang, J. F. Chen, C. Liu, S. Zhou, M. M. T. Loy, G. K. L. Wong, and S. Du. A dark-line two-dimensional magneto-optical trap of ^{85}Rb atoms with high optical depth. *Review of Scientific Instruments*, 83(7):073102, July 2012. [165](#), [231](#)

-
- [402] K. Dieckmann, R. Spreew, M. Weidemüller, and J. Walraven. Two-dimensional magneto-optical trap as a source of slow atoms. *Physical Review A*, 58(5):3891–3895, November 1998. 165
- [403] B. Wang, Y. Han, J. Xiao, X. Yang, C. Zhang, H. Wang, M. Xiao, and K. Peng. Preparation and determination of spin-polarized states in multi-Zeeman-sublevel atoms. *Physical Review A*, 75(5):051801(R), May 2007. 166, 217, 231
- [404] K. E. Miller, D. Krause Jr., and L. R. Hunter. Precise measurement of the Stark shift of the rubidium and potassium D1 lines. *Physical Review A*, 49(6):5128–5130, 1994. 167, 241, 243
- [405] J. J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, rev. edition, 1994. 170
- [406] J. Nunn, U. Dorner, P. Michelberger, K. Reim, K. Lee, N. Langford, I. Walmsley, and D. Jaksch. Quantum memory in an optical lattice. *Physical Review A*, 82(2):022327, August 2010. 173
- [407] T. Barter. *Measuring Gravity with a Bose-Einstein Condensate*. Honours thesis, The Australian National University, 2010. 183, 185
- [408] T. P. Dinneen, C. D. Wallace, and P. L. Gould. Narrow linewidth, highly stable, tunable diode laser system. *Optics Communications*, 92(4-6):277–282, September 1992. 187, 214
- [409] G. D. McDonald. *Detecting Atomic Shot Noise On Ultra-cold Atom Clouds*. Honours thesis, The Australian National University, 2009. 189
- [410] S. Burks. *Towards A Quantum Memory For Non-Classical Light With Cold Atomic Ensembles*. PhD thesis, L’Université Pierre et Marie Curie, 2010. 202, 204
- [411] S. Burks, J. Ortalo, A. Chiummo, X. Jia, F. Villa, A. Bramati, J. Laurat, and E. Giacobino. Vacuum squeezed light for atomic memories at the D2 cesium line. *Optic Express*, 17(5):3777–3781, 2009. 202
- [412] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen. Photon storage in Λ -type optically dense atomic media. II. Free-space model. *Physical Review A*, 76(3):033805, September 2007. 207
- [413] L. Giner, L. Veissier, B. M. Sparkes, A. S. Sheremet, A. Nicolas, O. S. Mishina, M. Scherman, S. Burks, I. Shomroni, D. V. Kupriyanov, P. K. Lam, E. Giacobino, and J. Laurat. Experimental Investigation of the Transition between Autler-Townes Splitting and Electromagnetically-Induced Transparency. *Physical Review A*, 013823(1):2013, 2012. 209, 210
- [414] W. R. Kelly, Z. Dutton, J. Schlafer, B. Mookerji, and T. A. Ohki. Direct Observation of Coherent Population Trapping in a Superconducting Artificial Atom. *Physical Review Letters*, 104(16):163601, April 2010. 208
- [415] Z.-H. Li, Y. Li, Y.-F. Dou, and J.-X. Zhang. The difference in noise property between the Autler-Townes splitting medium and the electromagnetically induced transparent medium. *Chinese Physics B*, 21(3):034204, March 2012. 208

- [416] A. A. Abdumalikov, O. Astafiev, A. M. Zagoskin, Y. A. Pashkin, Y. Nakamura, and J. S. Tsai. Electromagnetically Induced Transparency on a Single Artificial Atom. *Physical Review Letters*, 104(19):193601, May 2010. 208
- [417] A. H. Safavi-Naeini, T. P. Mayer Alegre, J. Chan, M. Eichenfield, M. Winger, Q. Lin, J. T. Hill, D. E. Chang, and O. Painter. Electromagnetically induced transparency and slow light with optomechanics. *Nature*, 472(7341):69–73, April 2011. 208
- [418] H. Tanji-Suzuki, W. Chen, R. Landig, J. Simon, and V. Vuletić. Vacuum-induced transparency. *Science*, 333(6047):1266–1269, September 2011. 208
- [419] P. Anisimov and O. Kocharovskaya. Decaying-dressed-state analysis of a coherently driven three-level Λ system. *Journal of Modern Optics*, 55(19-20):3159–3171, November 2008. 208
- [420] T. Y. Abi-Salloum. Electromagnetically induced transparency and Autler-Townes splitting: Two similar but distinct phenomena in two categories of three-level atomic systems. *Physical Review A*, 81(5):053836, May 2010. 208
- [421] P. M. Anisimov, J. P. Dowling, and B. C. Sanders. Objectively Discerning Autler-Townes Splitting from Electromagnetically Induced Transparency. *Physical Review Letters*, 107(16):163604, October 2011. 208, 210
- [422] H. Akaike. Information theory and an extension of the maximum likelihood principle. *International Symposium on Information Theory, 2nd, Tsahkadsor, Armenian SSR*, pages 267–281, 1973. 208
- [423] O. S. Mishina, M. Scherman, P. Lombardi, J. Ortalo, D. Felinto, A. S. Sheremet, A. Bramati, D. V. Kupriyanov, J. Laurat, and E. Giacobino. Electromagnetically induced transparency in an inhomogeneously broadened Λ transition with multiple excited levels. *Physical Review A*, 83(5):053809, May 2011. 210
- [424] M. Scherman, O. S. Mishina, P. Lombardi, E. Giacobino, and J. Laurat. Enhancing electromagnetically-induced transparency in a multilevel broadened medium. *Optics Express*, 20(4):4346–4351, February 2012. 210
- [425] M. T. DePue, S. L. Winoto, D. J. Han, and D. S. Weiss. Transient compression of a MOT and high intensity fluorescent imaging of optically thick clouds of atoms. *Optics Communications*, 180(13):73–79, 2000. 215
- [426] N. P. Robins, B. M. Sparkes, and Dedman C. High speed switching of magnetic field gradients. *In preperation*. 224
- [427] T. Horrom, A. Lezama, S. Balik, M. D. Havey, and E. E. Mikhailov. Quadrature noise in light propagating through a cold 87Rb atomic gas. *Journal of Modern Optics*, 58(21):1936–1941, December 2011. 232
- [428] P. Kolchin, S. Du, C. Belthangady, G. Y. Yin, and S. E. Harris. Generation of Narrow-Bandwidth Paired Photons: Use of a Single Driving Laser. *Physical Review Letters*, 97(11):113602, September 2006. 232
- [429] T. J. Weinhold. Personal communication, 2013. 238

-
- [430] S. Armstrong. Personal communication, 2013. 238
- [431] G. T. Campbell, O. Pinel, M. Hosseini, B. C. Buchler, and P. K. Lam. Arbitrary unitary transformations on optical states using a quantum memory. *In preperation*. 239
- [432] D. A. Steck. Cesium D line data. *available online at <http://steck.us/alkalidata> (revision 2.0.4, 23 December 2010)*. 241
- [433] B. M. Sparkes, H. M. Chrzanowski, D. P. Parrain, B. C. Buchler, P. K. Lam, and T. Symul. A scalable, self-analyzing digital locking system for use on quantum optics experiments. *Review of Scientific Instruments*, 82(7):075113, July 2011. 246, 270
- [434] Analog Devices Homepage. <http://www.analog.com/en.index.html>. 248, 261
- [435] National Instruments Homepage. <http://www.ni.com/>. 261
- [436] E. D. Black. An introduction to PoundDreverHall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, 2001. 271