

# High Performance Solid State Quantum Memory

Morgan P. Hedges

A thesis submitted for the degree of  
Doctor of Philosophy (Physics)  
Australian National University

August, 2011



Michael P. Hughes

A book about the history of the  
University of Queensland  
and its role in the development  
of Queensland.

1994

---

# Statement of Authorship

---

Except where due reference is made in the text, the material presented in this thesis is entirely my own work.



---

Morgan P. Hedges

August, 2011

# Statement of Antitrust

Statement of Antitrust

*[Faint signature]*

---

# Acknowledgements

---

When I first submitted this thesis, I did not feel positive enough to write acknowledgements. The thesis represented a goal which fell far short of justifying the journey, something which I could see with painful clarity throughout most of the 18 month writing process. I felt most of it better forgotten. Yet with 8 months past since the submission, the concept is easier to stomach. It's now obvious that the things one wants to forget are usually those that need acknowledging. I'll start by acknowledging that fact! Now, some people who have indirectly or directly influenced this document.

My family. Mum, Dad, and Mark. For large stretches of time, my studies caused me to lose sight of how important you are to me. Often it seemed more important to beat my head against some particular problem in the lab than to visit or call. So many stupid times it was, if I am honest, a feeling of insecurity about my progress and fear of disappointing you, that stopped me. But I really do know how silly that is. At my most miserable it has been a tremendous safety blanket to know that you are always there, totally accepting and ready to catch me no matter how badly I fall. I wish I could have been brave enough to really talk the times that I might have helped in return. Mark: I found it much easier to open up to you, for reasons that perhaps you know best. I am certain that you don't realise how important our late night calls have been to rescue me from some quite black moods. Please don't stop answering.

Erin. Without you, I am not sure how this little venture would have gone, but I know it would have been lesser. Perhaps I would have fallen over the edge completely, or perhaps I would have passed through with a short PhD and very little else. But you helped me explore a much wider landscape. Thanks to you I've seen some very exciting hills that I'll spend the rest of my life discovering, and I can navigate better around the holes. I am improved for learning what you've taught. You have my utmost respect and appreciation.

My colleagues: one could not ask for better. Wonderful, genuine, highly intelligent folk. I must give particular mention to Sarah, who accompanied me the whole way through. The experience has been much richer, academically and otherwise, thanks to you. John whose intelligent patience will crack any problem, and Rose who can name and fix the problems that have long been ignored from embarrassment. Thank you both. Sara and Kate, I wish I could have had more time to be involved with your projects, and to get to know you better. You have certainly helped with a positive working environment. That leaves Milos and Grace: It has been a relatively brief experiment in supervision, which I have enjoyed immensely. I just hope you will recover from the experience! I would be honoured to work with you in the future.

Jevon and Neil, two fine examples of scientists who make the academic life look worth pursuing. You have been stabilising, motivating influences without whom things would have been much tougher. But I must also say: Jevon, without your thesis to haunt me, I'd be somewhat happier with mine!

Matt, officially the longest-term colleague I've ever had, and undoubtedly one of the finest scientists I have met. It's been a fascinating ride. You have taught me more than I can currently get my head around. Certainly the value of independence, the necessity of strong principles and the importance of believing in ones own ideas in science. I think you will understand if I say that you've led both by shining example and tragic anti-example. I will long remember and use both to guide my future steps. If only you could have taught me to publish regularly! There are many who would appreciate reading more of your work. You've brought out and dealt with the best and the worst in me. I hope you'll remember the best.

Finally, thanks to myself. For realising that this document will be read and appreciated largely only by myself, while maintaining enough self belief to stay the right side of the edge and eventually finish the beast. Sure, the experiences of stronger folk might make my own look petty. But there are also many intelligent and tough potential scientists, several of whom I know and respect, that have found the hurdles much taller. Whether by my own efforts or complete luck, with flying colors or by a whisker, I have finished. And that's worth acknowledging.

---

# Abstract

---

This thesis demonstrates an efficient quantum memory that for the first time preserves more quantum information than it loses. Beyond this, it describes how even higher efficiency may be realised without sacrificing the other performance characteristics that are required for a practical device.

A quantum memory is a device to store and recall the quantum state of light. Such a memory will be required for most applications of the emerging field of quantum information. The closest practical application is quantum repeaters for long distance encryption. A more distant target is quantum computing. Yet demonstrations of quantum memory to date fall far short of the performance required for repeaters, and further short of that for quantum computing. For these applications, high performance will be required in terms of 3 parameters: long storage time, high efficiency, and a large mode-capacity.

Prior to this work, the storage time was the only parameter to have been demonstrated near the required level[1], and that only in a classical regime. The primary result of this thesis is to demonstrate a memory with an efficiency close the required level (70% is measured), and confirm its quantum nature with noise measurements. This represents the first quantum memory with efficiency greater than 50% and the first to operate unconditionally above the “no-cloning limit”. The demonstration used the same storage medium as for the previous long-term storage experiment ( $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$  (PrYSO)), and could in principle reach the same long storage times. The technique used is an optical “gradient echo”, an optical equivalent to the long studied gradient echo technique used in Nuclear Magnetic Resonance (NMR).

Aside from its practicality for demonstrations, a gradient-echo also offers a powerful handle for manipulating stored light. However it is more complex than related memory techniques because it employs a spatially-varying medium. This becomes crucial when the optical depth (and memory efficiency) is increased; propagation effects mean the light is stored very differently than in other techniques. A simple linear-response model is used to gain insight into this high optical depth regime. As a result, some new properties are identified which may increase the efficiency in practical circumstances and improve flexibility. Of particular note, it's found that the space-dependence can be used to hide one narrow transition with another. This allows the somewhat unintuitive notion that a light pulse can fully traverse a medium without exciting a transition that it is resonant with.

The spectral-hiding property has promise to overcome the problem with simultaneously achieving all three parameters required for quantum memories. The problem is related to the preparation of available rare-earth ion doped materials. To achieve high efficiency

and storage times, these require spectral hole-burning whose bandwidth is limited by the splittings of spin levels. Methods exist to create high bandwidth delays lines with holeburning[2], but they cannot be efficient or long-lived. By utilising the spectral-hiding property of the gradient echo, methods to overcome this limitation are described for the first time.



---

# Contents

---

Statement of Authorship	iii
Overview	ix
Glossary	xvii
<b>1 Background</b>	<b>3</b>
1.1 Field of Quantum Information . . . . .	3
1.1.1 Quantum Bits . . . . .	5
1.1.2 Systems of Qubits . . . . .	6
1.1.3 Entanglement and Entanglement Swapping . . . . .	6
1.1.4 Quantum Information with Continuous Variables . . . . .	8
1.1.5 Quantum Networking . . . . .	8
1.1.6 Quantum Repeaters . . . . .	10
1.2 Quantum Optics . . . . .	12
1.2.1 Modes of the Electromagnetic Field . . . . .	12
1.2.2 Quantization of a Single Mode . . . . .	14
1.2.3 Quadrature Amplitudes . . . . .	15
1.2.4 Photon Number States . . . . .	15
1.2.5 Coherent and Gaussian States . . . . .	16
1.2.6 Beam-splitters and Amplifiers . . . . .	17
1.2.7 Counting Photons . . . . .	18
1.2.8 Measuring Amplitude . . . . .	19
1.3 Atom-light Interactions . . . . .	28
1.3.1 Approximate Picture for Light-Ensemble Interaction . . . . .	29
1.3.2 Damped, Driven, Dipole Oscillator . . . . .	30
1.3.3 Re-radiation by Damped Dipoles . . . . .	32
1.3.4 Optical Bloch Equations . . . . .	33
1.3.5 Bloch Sphere Dynamics . . . . .	35
1.3.6 Maxwell-Bloch Equations . . . . .	38
1.3.7 Spectral Response, Absorption and Dispersion . . . . .	39
1.3.8 Slow Light . . . . .	40
1.3.9 Population Inversion and Amplification . . . . .	42
1.3.10 Impulse Response . . . . .	42
1.3.11 Time-dependant Response: Echoes . . . . .	44

---

1.3.12	Transferring Coherence and Changing the Emission Direction . . . .	46
<b>2</b>	<b>Coherent and Quantum Memories in Atomic Ensembles</b>	<b>49</b>
2.1	Classical Memories, Quantum Cloners, Amplifiers, and Measurement . . . .	50
2.2	General Ensemble-based Coherent memory . . . . .	53
2.2.1	Total Efficiency . . . . .	54
2.2.2	Coupling Efficiency . . . . .	54
2.2.3	Bandwidth . . . . .	55
2.2.4	Decay Time and Shelving Time . . . . .	55
2.2.5	Time-bandwidth Product and Mode Capacity . . . . .	55
2.2.6	Background and Reabsorption . . . . .	56
2.2.7	Noise . . . . .	56
2.3	Linear, Coherent Memory Techniques . . . . .	57
2.3.1	Electromagnetically Induced Transparency and Slow Light . . . . .	57
2.3.2	2-pulse Photon Echo . . . . .	59
2.3.3	3-pulse Photon-Echo and Atomic Frequency Combs . . . . .	60
2.3.4	Controlled Reversible Inhomogeneous Broadening and Optical Gradient Echoes . . . . .	60
2.4	Comparison with Finite Optical Depth . . . . .	61
2.5	Heralded Memories and DLCZ . . . . .	63
2.6	Low-finesse Cavities . . . . .	65
2.7	Quantum Benchmarks and Performance Measures . . . . .	65
2.7.1	A Hypothetical Pre-amplifier . . . . .	66
2.7.2	A Pre-amplifier and Amplitude Noise . . . . .	67
2.7.3	Quantum Fidelity for Memories . . . . .	68
2.7.4	Noise Corrected Efficiency . . . . .	69
2.7.5	Conditional Measurements and Fringe-visibility . . . . .	70
2.7.6	State and Process Tomography . . . . .	71
2.8	Memories in Quantum Repeaters . . . . .	71
2.8.1	Generation and Swapping . . . . .	73
2.8.2	Modes Shuffled or in Parallel . . . . .	73
2.8.3	Single-Photon Entanglement and Partial Readout . . . . .	73
2.9	Meeting the Requirements . . . . .	74
<b>3</b>	<b>Rare-earth Ions and Spectral Hole-burning for Memories</b>	<b>75</b>
3.1	Why Rare-earth Ions . . . . .	76
3.2	Spectra of 4f-4f Transitions in Crystals . . . . .	77
3.2.1	A Typical Spectrum of a Rare-earth Doped Crystal . . . . .	78
3.2.2	Free-ion Approximation . . . . .	78
3.2.3	Crystal Field Levels . . . . .	81
3.2.4	Hyperfine Structure in Zero Field . . . . .	83

---

3.2.5	Zeeman Structure . . . . .	83
3.2.6	Lifetimes . . . . .	84
3.2.7	Transition Strength . . . . .	84
3.2.8	Homogeneous Broadening . . . . .	85
3.2.9	Inhomogeneous Broadening . . . . .	86
3.2.10	Reducing Dephasing . . . . .	87
3.3	Tailored Spectra using Metastable Hole-burning . . . . .	87
3.3.1	Population Density and Absorption Profiles . . . . .	89
3.3.2	Power Broadening . . . . .	90
3.3.3	Burn-back and Population Shuffling . . . . .	90
3.3.4	Coherent Optical Transients for Hole-burning . . . . .	90
3.3.5	Competing Pumping . . . . .	91
3.3.6	Directly Addressing Spin Transitions (Using RF) . . . . .	92
3.3.7	Limitations of Hole-burning . . . . .	92
3.4	Specific Materials for Quantum Memory . . . . .	94
3.4.1	Target Material Parameters . . . . .	94
3.4.2	High-Coherence Transitions . . . . .	95
3.4.3	Future Materials . . . . .	98
3.4.4	Further details of PrYSO . . . . .	99
3.5	Alternatives to Metastable Hole-burning . . . . .	100
3.5.1	Persistent Hole-burning . . . . .	100
3.5.2	Laser Induced transition . . . . .	102
3.5.3	Strain-free Crystals . . . . .	102
3.5.4	Variations of a 2-pulse Photon Echo . . . . .	102
<b>4</b>	<b>Optical Gradient-Echo - Theory and Modelling</b>	<b>105</b>
4.1	Gradient-Echo Model . . . . .	106
4.1.1	Units and Simulation Parameters . . . . .	107
4.2	Unmodified Optical Gradient-Echoes . . . . .	107
4.2.1	Space-dependent Response in a Gradient Echo . . . . .	108
4.2.2	A Time-domain Memory . . . . .	111
4.2.3	A Frequency Domain Memory . . . . .	114
4.2.4	A Polariton in Time and k-space . . . . .	116
4.3	Gradient-Echo with Inhomogeneously Broadened Transitions . . . . .	116
4.3.1	Inhomogeneous Lorentzian . . . . .	117
4.3.2	Inhomogeneous Top-hat . . . . .	117
4.3.3	Hidden Transitions . . . . .	120
4.4	General Detuning Fields and Manipulating Light . . . . .	122
4.4.1	Shelving and Suppressing . . . . .	122
4.4.2	Field Switching Speed . . . . .	123
4.4.3	Manipulating the Output Mode . . . . .	124

---

4.4.4	Repeated Retrieval to Increase Read Efficiency . . . . .	124
4.4.5	Non-monotonic Fields and Multi-stage Memories . . . . .	125
4.4.6	Pulse Sequencing and Echo Suppression . . . . .	128
4.5	Improving Efficiency . . . . .	128
4.5.1	An Efficient Initial Feature . . . . .	130
4.5.2	An Efficient Pulse Shape . . . . .	130
<b>5</b>	<b>Realisation of an Efficient Gradient Echo</b>	<b>135</b>
5.1	Technical Considerations . . . . .	136
5.1.1	General Experiment Setup . . . . .	136
5.1.2	Laser Stability . . . . .	139
5.1.3	Sample Stability . . . . .	140
5.1.4	Interferometer Stability . . . . .	140
5.1.5	Detectors and Balancing . . . . .	142
5.2	Hole-filling in PrYSO . . . . .	144
5.2.1	Crystal Choice . . . . .	147
5.3	Creating and Measuring a High-contrast Anti-hole . . . . .	147
5.3.1	Making a Transmissive Spectral Trench . . . . .	147
5.3.2	Using All Possible Transitions . . . . .	150
5.3.3	Avoiding Propagation Effects by Preparing from the Side . . . . .	151
5.3.4	Robust Preparation by Competing Processes . . . . .	152
5.3.5	Cleaning up the Background . . . . .	154
5.3.6	Estimating an Optically Thick Feature . . . . .	155
5.4	Monotonic Stark Shift in PrYSO . . . . .	156
5.4.1	Spatial Stark Shift from a Simple Quadrupole . . . . .	157
5.4.2	Stark Shift due to a Rotated Block-Quadrupole . . . . .	159
5.5	Efficient Gradient Echoes . . . . .	162
5.5.1	Preparation Sequence . . . . .	162
5.5.2	Efficiency Measurements . . . . .	164
5.5.3	Predicted Improvements . . . . .	166
5.6	Quantum noise measurements . . . . .	167
5.6.1	Quadrature Measurements . . . . .	168
5.6.2	Maximum Likelihood Estimation of Covariance . . . . .	170
5.6.3	Attenuation and Additional Noise . . . . .	172
5.6.4	Comparison with Benchmarks . . . . .	173
5.6.5	Added Phase Noise . . . . .	175
5.6.6	Phase Modulation of the Temporal Mode . . . . .	176
5.7	Other Demonstrations . . . . .	177
5.7.1	Multiple Stark Groups . . . . .	177
5.7.2	The Edge Effect . . . . .	178

---

<b>6</b>	<b>Prospects for Practical Quantum Memory using Rare Earths</b>	<b>181</b>
6.1	Improvements to the Current Experiment . . . . .	182
6.1.1	Improving Efficiency and Time-Bandwidth Product . . . . .	182
6.1.2	Long-term Shelving . . . . .	182
6.2	Techniques for a Real Device . . . . .	183
6.2.1	Highly Parallel Hole-burning Memories . . . . .	184
6.2.2	Quiet Optical Rephasing . . . . .	185
6.2.3	Raman Feature with a Spin-Polarised Sample . . . . .	186
6.2.4	Hyperfine-Resolved Stoichiometric Crystals . . . . .	188
6.3	Memories With and As a Quantum Light Source . . . . .	190
6.4	Summary . . . . .	191
	<b>Bibliography</b>	<b>193</b>

---

101	1.1. Introduction	101
102	1.2. The structure of the book	102
103	1.3. Acknowledgements	103
104	1.4. References	104
105	2. The structure of the book	105
106	2.1. The structure of the book	106
107	2.2. The structure of the book	107
108	2.3. The structure of the book	108
109	2.4. The structure of the book	109
110	2.5. The structure of the book	110
111	2.6. The structure of the book	111
112	2.7. The structure of the book	112
113	2.8. The structure of the book	113
114	2.9. The structure of the book	114
115	2.10. The structure of the book	115
116	2.11. The structure of the book	116
117	2.12. The structure of the book	117
118	2.13. The structure of the book	118
119	2.14. The structure of the book	119
120	2.15. The structure of the book	120
121	2.16. The structure of the book	121
122	2.17. The structure of the book	122
123	2.18. The structure of the book	123
124	2.19. The structure of the book	124
125	2.20. The structure of the book	125
126	2.21. The structure of the book	126
127	2.22. The structure of the book	127
128	2.23. The structure of the book	128
129	2.24. The structure of the book	129
130	2.25. The structure of the book	130
131	2.26. The structure of the book	131
132	2.27. The structure of the book	132
133	2.28. The structure of the book	133
134	2.29. The structure of the book	134
135	2.30. The structure of the book	135
136	2.31. The structure of the book	136
137	2.32. The structure of the book	137
138	2.33. The structure of the book	138
139	2.34. The structure of the book	139
140	2.35. The structure of the book	140
141	2.36. The structure of the book	141
142	2.37. The structure of the book	142
143	2.38. The structure of the book	143
144	2.39. The structure of the book	144
145	2.40. The structure of the book	145
146	2.41. The structure of the book	146
147	2.42. The structure of the book	147
148	2.43. The structure of the book	148
149	2.44. The structure of the book	149
150	2.45. The structure of the book	150
151	2.46. The structure of the book	151
152	2.47. The structure of the book	152
153	2.48. The structure of the book	153
154	2.49. The structure of the book	154
155	2.50. The structure of the book	155
156	2.51. The structure of the book	156
157	2.52. The structure of the book	157
158	2.53. The structure of the book	158
159	2.54. The structure of the book	159
160	2.55. The structure of the book	160
161	2.56. The structure of the book	161
162	2.57. The structure of the book	162
163	2.58. The structure of the book	163
164	2.59. The structure of the book	164
165	2.60. The structure of the book	165
166	2.61. The structure of the book	166
167	2.62. The structure of the book	167
168	2.63. The structure of the book	168
169	2.64. The structure of the book	169
170	2.65. The structure of the book	170
171	2.66. The structure of the book	171
172	2.67. The structure of the book	172
173	2.68. The structure of the book	173
174	2.69. The structure of the book	174
175	2.70. The structure of the book	175
176	2.71. The structure of the book	176
177	2.72. The structure of the book	177
178	2.73. The structure of the book	178
179	2.74. The structure of the book	179
180	2.75. The structure of the book	180
181	2.76. The structure of the book	181
182	2.77. The structure of the book	182
183	2.78. The structure of the book	183
184	2.79. The structure of the book	184
185	2.80. The structure of the book	185
186	2.81. The structure of the book	186
187	2.82. The structure of the book	187
188	2.83. The structure of the book	188
189	2.84. The structure of the book	189
190	2.85. The structure of the book	190
191	2.86. The structure of the book	191
192	2.87. The structure of the book	192
193	2.88. The structure of the book	193
194	2.89. The structure of the book	194
195	2.90. The structure of the book	195
196	2.91. The structure of the book	196
197	2.92. The structure of the book	197
198	2.93. The structure of the book	198
199	2.94. The structure of the book	199
200	2.95. The structure of the book	200
201	2.96. The structure of the book	201
202	2.97. The structure of the book	202
203	2.98. The structure of the book	203
204	2.99. The structure of the book	204
205	2.100. The structure of the book	205

---

# Overview

---

The principle goals of this thesis are to realise an efficient quantum memory, and to explain how this work may be further developed to realise a memory with practical performance levels in all aspects. To cover the necessary material, a topical structure is followed from the general to the specific. The experiments presented are relatively narrow in scope and use similar techniques. They are thus grouped and given toward the end (in chapter 5) to avoid repeating details.

The topical approach is taken for two main reasons. Firstly it may help to ease the burden on readers from different areas. This work sits somewhere between the usual considerations of workers in quantum optics and in solid-state spectroscopy. Hopefully the topical progression allows the reader to focus on the parts of their interest. The second reason is to aid a coherent discussion of the broader challenges and directions of the field<sup>1</sup>.

A synopsis of each chapter follows.

Chapter 1 is Background. It aims to cover the bare necessary elements in three areas. The field of quantum information motivates the study of quantum memory, and thus the present work. Quantum optics is necessary to describe the light that is stored by a memory, and how to measure it. Finally a knowledge of atom-light interactions is necessary to understand memory implementations.

Chapter 2 is Quantum and Coherent Memories in Atomic Ensembles. Common techniques for implementing memories using ensembles of absorbers are compared. A definition is given for what it means for a memory to be ‘quantum’, and benchmarks are derived to determine this in an experiment. The requirements of a quantum memory for implementing a useful quantum repeater are discussed. The main requirements are found to be usable optical depth, and a long enough coherence time.

Chapter 3 is Rare-earth Ions for Memories. Rare-earth ions in the solid state meet the combined requirements of high optical depth and long coherence time very strongly. They are also very flexible and able to implement any ensemble based technique thanks to spectral hole-burning. General spectroscopy and techniques for optimising hole-burning are covered, with the purpose of implementing memories. Memories implemented using hole-burning are found to be limited, particularly in bandwidth, by material properties. A discussion of available materials indicates that while suitable for near-term demands, present techniques will be limited for applications in the long run. We also identify PrYSO as the most likely material for a high-efficiency demonstration.

---

<sup>1</sup>The third reason is that I didn’t realise how long it would take

Chapter 4 is Optical Gradient-Echo - Theory and Modelling. This technique is chosen in the first place as the most practical for demonstrating a highly efficient memory. This chapter studies the technique in detail, while attempting to give an intuitive picture of its quirks due to spatial dependence. It is apparent after this study that beyond a demonstration, the technique is also highly promising also in the long term. It is found that the gradient echo has lower requirements than originally thought, has the flexibility to nearly arbitrarily manipulate the light mode to improve repeater applications, and ultimately has the ability to ‘hide’ spectral structure. This may provide a key to removing bandwidth limitations of Rare-earth ion based memories discussed in chapter 3. The unusual propagation properties of the memory are illustrated using numerical modelling. Experimental requirements and optimisation are also discussed.

Chapter 5 is Realisation of an Efficient Gradient-Echo. A gradient echo is implemented using spectral hole-burning in PrYSO. An extremely absorptive optical feature is created and characterised using spectral hole-burning. The dielectric and Stark-shift related properties of PrYSO are investigated and modelled. High efficiency gradient-echoes are then demonstrated and compared with theory. Finally measurements of quantum noise are made and the memory is found to be ‘quantum’.

Chapter 6 is Practical Quantum Memory in Rare Earths. It offers a brief summary of what has been learned from this work and identifies new directions. Bringing together items discussed in previous chapters, combined strategies for quantum memories in solids containing Rare-earth ion (REI)s are discussed. The emphasis is to make full, efficient use of the absorbers in a REI material for real applications.



---

# Glossary

---

- 2PE** Two-Pulse Photon Echo. 54, 57, 59, 61, 100, 177, 182
- AFC** Atomic Frequency Comb. 54, 58, 60, 61, 91, 123, 127, 181
- AOM** Acousto-Optic Modulator. 130, 131, 136, 159
- BEC** Bose-Einstein Condensate. 71
- COT** Coherent Optical Transient. 88
- CPO** Coherent Population Oscillation. 57
- CRIB** Controlled Reversible Inhomogenous Broadening. 54, 55, 58–61, 93, 101, 127, 180–182
- DDC** Dynamic Decoherence Control. 85, 96, 97, 175
- DLCZ** DLCZ. 61, 62, 69, 181, 182
- EIT** Electromagnetically Induced Transparency. 54, 55, 57, 59, 97
- ErYSO** Er<sup>3+</sup>:Y<sub>2</sub>SiO<sub>5</sub>. 94, 179
- ESR** Electron Spin Resonance. 90
- EuYSO** Eu<sup>3+</sup>:Y<sub>2</sub>SiO<sub>5</sub>. 95–97, 174, 175
- FID** Free-induction Decay. 27, 35, 42, 57, 115, 118, 160
- FWHM** full-width at half-maximum. 52, 53, 55, 125
- ISD** Instantaneous Spectral Diffusion. 84
- LO** Local Oscillator. 19–24, 130, 131
- MBE** Maxwell-Bloch Equations. 27, 35, 37, 38
- NMR** Nuclear Magnetic Resonance. 58, 85, 90
- OBE** Optical-Bloch Equations. 27, 30–32, 35, 38

- 
- PrYSO**  $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ . 96–98, 129, 130, 133, 137, 139, 140, 142–145, 148, 149, 153, 154, 174–176, 179
- QKD** Quantum Key Distribution. 8
- QND** Quantum Non-Demolition. 61
- QST** Quantum State Tomography. 67
- RASE** Rephased Amplified Spontaneous Emission. 181, 182
- REI** Rare-earth ion. 59, 71, 73–76, 79, 85, 92, 95, 98, 101, 173, 176, 181, 183
- RWA** Rotating Wave Approximation. 29
- SPD** Single Photon Detector. 16, 17, 62, 68
- SPDC** Spontaneous Parametric Down Conversion. 181
- SVEA** Slowly Varying Envelope Approximation. 29, 37, 40, 42
- TBP** Time-Bandwidth Product. 53, 54, 57, 61, 92, 94, 95, 173–176, 183
- TDOM** Time-domain Optical Memory. 92, 93, 96
- YAG**  $\text{Y}_3\text{Al}_5\text{O}_{12}$ . 93, 94
- YALO**  $\text{YAlO}_3$ . 81
- YLF**  $\text{LiYF}_4$ . 180
- YSO**  $\text{Y}_2\text{SiO}_5$ . 93–95, 97, 98, 121, 174
- ZFOZ** Zero First-Order Zeeman. 85, 95–97, 175



---

# Background

---

A “quantum memory for light” is a simple device. Its purpose is to interact with a pulse of light carrying quantum information to control what time it arrives. A quantum memory cannot use measurement and requires a high efficiency for a single reason: quantum information cannot be copied, and thus any record or trace of the pulse that remains after the pulse has passed represents information lost.

The purpose of this thesis is to present a few steps on the long path toward practical quantum memory technology, particularly toward memory efficiency. The purpose of this chapter is to give a brief background of the why, the what, and the how for building a quantum memory. The why is the application field of quantum information, the what is the quantum properties of light, and the how is the interaction of light with ensembles of atoms.

The challenges of implementing efficient quantum memory are parallel with the broader field of quantum information: interacting with and controlling a system which is otherwise totally isolated from the noisy, classical world. A quantum memory is thus not only an enabling technology for quantum computing and networking applications; it’s strongly related to all aspects. A broad view of the field of quantum information is given in the first section of this chapter

What a quantum memory does is preserve the quantum state of light. The field concerning these properties is quantum optics. The second section of this chapter is devoted to describing relevant specifics- the quantum nature of light, how to measure it, and how it changes with loss.

The quantum memories we study are ensembles of atomic absorbers, which interact with, store, and (the tricky part) recall incident light. The last section of the chapter gives an overview of the interaction of light with idealised atomic ensembles. We emphasise a simple semi-classical description of two level atoms, as this model is both intuitive and adequate for the majority of this work.

## 1.1 Field of Quantum Information

In the 20th century, a huge research effort into information processing occurred, as the power of computers manipulating this information increased exponentially. At nearly the

same time, so did the complexity of the networks connecting them. Apart from changing the way we keep up on the latest gossip, this has drastically changed the way we use and store information. Today the Internet Archive[3] stores over 150 billion web pages, while the Google Books project has already created digital copies of 15 million books[4] out of an estimated 130 million ever written[5]. Even Wikipedia[6] currently lists 4 million articles, compared to the estimated 500,000 scrolls that an earlier repository of human knowledge, the Great Library of Alexandria, held at its peak[7]. This is all available within a good fraction of the world's homes, and even an increasing number of pockets through smart phones.

For the moment, the emerging field of quantum information is not set to challenge current technology in these applications. Quantum information is, however, an extension to the 'classical' information used by today's computers (and by Wikipedia). Such classical information has as a basis the smallest unit of logic, the 'bit'. A comparison with classical information and the classical bit can help us know what quantum information is and isn't.

Variations on binary bits, extending it to values other than 0 and 1, have also been considered. In one extreme, the field of analog computing where an analog 'bit' can take a continuum of values promised a huge increase in information density. Unfortunately however, such a scheme is just as limited due to the increasing complexity required to account for errors. In terms of the rough number of computations required to solve a problem, digital computers based on binary were thought to be at least as good as alternatives. The famous Church-Turing thesis[8] stated (roughly) that a Turing machine (of which a digital computer is a simple example) can compute anything for which an *effective algorithm* exists. A stronger interpretation of the Thesis has often been taken, that anything computable is computable on a Turing machine<sup>1</sup>. While the stronger interpretation holds for common notions of computing, there are certainly exceptions. In particular, it was realised in the 70s that some problems could be solved more quickly using a stochastic algorithm, that is by introducing some randomness. In some sense, complete certainty about the answer is given up in exchange for (potentially) much faster run-times.

With the benefit of hindsight, and a knowledge of quantum mechanics, it is perhaps easy to see a further exception: it is now well known that quantum systems cannot be simulated efficiently using classical logic, but such a system can certainly simulate itself. This was first noted by Feynman in 1982[10]. Deutsch in 1985[11] took the idea further, describing a universal computer based on two level quantum systems of 'qubits'. Such a computer can in principle simulate both quantum and classical physics (and logic) as or more 'efficiently' than a computer using just 0s and 1s.

A real, widely useful and applicable quantum computer is yet to be demonstrated, and there are to date only a handful of algorithms known to perform better on a quantum computer. At least 2 of these algorithms though —Shor's factoring into prime numbers and Grover's database search— are of particular relevance to present trends. As quantum

---

<sup>1</sup>An interesting discussion is given by Goldin: "The Church-Turing Thesis: Breaking the Myth"[9]




Figure 1.1: Bloch-sphere representation of a qubit. A single (pure) qubit is represented by a point on the surface, while an ensemble is in general represented by a point within the sphere. A classical probabilistic bit may be represented as point on the  $w$  axis.

computers become available, the pace of algorithm development should naturally increase.

There are also other uses for this new type of logic. Perhaps most attractive for a physicist, quantum computing promises the efficient simulation of quantum systems. The most immediate application though is using quantum information as a secure communication tool. That is, building quantum networks.

### 1.1.1 Quantum Bits

The qubit is so named as the quantum analogue to the classical ‘bit’. It is the smallest unit of quantum information. Any two-level quantum system may be thought of as a qubit, a regular example of which is a spin-1/2 particle (e.g. a nuclear spin) [12]. Systems with 3 and more levels (qutrits, qudits) have also been considered, but it is unclear whether there is a benefit to outweigh the added complexity.

Any qubit may be represented by a linear combination of two states  $|\psi\rangle = c_1|0\rangle + c_2|1\rangle$ , where of course  $c_1$  and  $c_2$  are complex numbers. Neglecting the global phase factor and noting the normalization restriction that  $|c_1|^2 + |c_2|^2 = 1$ , we see that a single qubit may be represented by two real numbers. A qubit with some classical uncertainty, or alternatively an ensemble of qubits, may be represented by a density matrix with three free parameters.

$$\rho = \begin{pmatrix} a & b_1 - ib_2 \\ b_1 + ib_2 & 1 - a \end{pmatrix} \quad (1.1)$$

The closest classical analog is a stochastic bit, which only needs a single parameter.

This may be written:

$$\rho_{class} = \begin{pmatrix} a & 0 \\ 0 & a - 1 \end{pmatrix} \quad (1.2)$$

On its own however, a single qubit can only yield the same *amount* of information as a classical bit. Only a single measurement can be made, and thus only one parameter accessed, before the state collapses. The difference is that the qubit gives a choice about which measurement to make, and thus which information to access. This property of choice already provides the basis for a scheme of secure communication[13].

### 1.1.2 Systems of Qubits

If we have  $N$  *independent* qubits instead of just 1, the number-of-numbers required to describe them is just  $3N$  (with no assumption that these are pure states). We can make  $N$  measurements to determine what the numbers are, so not much has changed.

When the qubits may be correlated, the information needed to describe the system scales very differently. There are  $2^N$  different possible states, each of requires two real numbers to specify it's complex amplitude. After including the probability normalising requirement,  $2^{N+1} - 1$  real numbers are required to specify the entire state. This is sometimes quoted as demonstrating the power of quantum computing, as a system of only  $\approx 300$  or so qubits would need more numbers than there are atoms in the universe to describe it.

It's worth noting, however, that classical bits may also be correlated. A system of  $N$  bits may be described as a linear combination of  $2^N$  different outcomes (each of which may be written  $(x_1 x_2 \dots x_N)$  for  $x_i = 0, 1$ ). Apart from the probability normalisation requirement, these outcomes may in principle be entirely independent<sup>2</sup> such that a number is required for each. The system requires  $2^N - 1$  real numbers to describe. In terms of the number-of-numbers needed to describe them, an arbitrary state of 301 stochastic bits is equivalent to 300 qubits.

Nevertheless, the strength and character of the correlations that may exist between qubits can be quite different from classical bits. These quantum correlations are given the name 'entanglement'.

### 1.1.3 Entanglement and Entanglement Swapping

The correlations that may exist between quantum states beyond that which exists for the probabilistic classical bits are referred to as entanglement. It is often considered the information resource in quantum information, as a way to quantify the ability of a system to process information in a non-classical way. It is generated, for example, when there

---

<sup>2</sup>Imagine a large collection of unlabelled boxes each containing a certain outcome configuration (say  $N$  compartments, each with a coin). If choosing a box at random, the probability of a particular outcome configuration depends only on the proportion of boxes prepared in that way.

is a truly random (in the quantum sense) outcome of a some event. Objects affected by the event are entangled until a record exists that determines what the outcome was (i.e. one of the outcomes was measured, or interacts with a larger system which could then in principle be measured to determine the outcome).

A common example is the decay of a single photon by some mechanism into two photons of lower energy, as in Spontaneous Parametric Down Conversion (SPDC). The properties of the two generated photons must be correlated to conserve energy and momentum, but the way the quantities are shared varies in a truly random fashion. When one photon interacts with a detector and is measured<sup>3</sup> the properties of the other may also be inferred. Before that time, the properties of the two photons are uncertain in an absolute sense. No information exists anywhere in the universe about how the properties are shared.

On the surface this is no different from a standard, "classical" correlation. Consider a common situation for those of us with matched left and right socks. Taking one of a pair in each hand, it is (in the opinion of the author) impossible to tell which is the left and which is the right until a measurement is made by attempting to wear one sock. At this point, information is gained also about the unmeasured sock. This is similar to the photons, but in this case it is merely the sock owner who is ignorant, not the universe.

So what's the difference? For one, socks don't have phase, and they don't have the Heisenberg uncertainty principle governing the measurement of conjugate observables. Only one kind of measurement can be made about the sock's handedness: one in the left-right basis. For a qubit though it is just as physical to measure in an up-down basis, or at some arbitrary angle. The photon qubits generated by processes like down conversion are correlated in all measurement bases. Measurements in these other bases yield additional information not available to merely correlated socks, and indeed not explainable by a theory where it is merely the experimenter who is ignorant, while the universe "knows" the measurement outcome ahead of time (so-called hidden-variable theories). The term 'entanglement' is used to denote the extra, uniquely quantum correlation.

More complete discussions are given in most textbooks (e.g. Nielsen and Chuang[14]). For the purposes of this thesis we merely need to know that entanglement exists, and that it can be 'swapped'[15]. Take two entangled pairs of quantum objects (e.g. photons)  $A$ - $B$  and  $C$ - $D$  which may be represented as  $|\psi\rangle_{ab}$  and  $|\psi\rangle_{cd}$ .  $A$  is entangled with  $B$ , and  $C$  with  $D$ . A joint measurement of  $B$  and  $C$  is performed; this reveals some combined property of the two, without revealing the state of the individual objects. This creates a quantum correlation (entanglement) between  $A$  and  $D$ , so the entanglement is 'swapped'. This principle is used widely in quantum information. An analogous process with correlated socks would be to take one sock from each of two identical pairs. If we find that we have taken a left and a right sock, we then know the remaining two also form a complete pair and are similarly correlated as the original pairs.

---

<sup>3</sup>Here "detector" is general. It may be any large system with which the photon interacts. The requirement is that the system can in principle be measured to determine the photon properties.



### 1.1.4 Quantum Information with Continuous Variables

For the purposes of quantum information, qubits may be quite directly realised using the discrete properties of the photon. Values of 0 and 1 can be encoded for example in the presence or not of a photon, its time of arrival, its polarisation, or the choice of path on a beam splitter. The correspondence is clear because the observables corresponding to a photon are discrete.

Particle like states are not the only choice for quantum information however. Many experiments have used states closer to the wave-like nature of light, with the variable of interest being the amplitude. This falls under the umbrella of continuous variable quantum information. Such continuous information may be considered analogously to analog classical computing. Alternatively, qubits may be mimicked in the same way as a continuous voltage can be used to convey digital information. An example of qubits realised using amplitude are ‘cat’ states, where the qubits (or qudits) are formed as a superposition of states of clearly distinguishable amplitude: that is something like  $c_1|-\alpha\rangle + c_2|\alpha\rangle$ , where  $\alpha$  is some mean amplitude.

One difference that often holds between the approaches is how loss is dealt with. For qubits based on single photons, 50% loss means 50% of the qubits are lost, but those that aren’t are unchanged. For continuous variables, loss means all pieces arrive, but each is degraded in quality. This is reflected in the detectors used to measure each.

Importantly for us however, the requirements for a quantum memory are the same, low loss and low noise. We are not so concerned with which approach, but in the practical parts of this work the available light source (i.e. a laser) corresponds to the continuous-variable regime. To fully characterise this light it’s amplitude must be determined and thus we use direct amplitude (i.e. continuous variable) measurements as realised by homodyne detectors (see section 1.2.8)<sup>4</sup>.

This is the natural choice. The main motivation for this choice is the light source we have available, which is a laser. To fully characterise these states, phase sensitive measurements must be made, so that the natural choice is homodyne detection. We note that interference measurements with photon counters (which were also readily available) could also have been used.

Continuous-variable quantum information is described extensively by Lloyd et al. [16], the review article by Braunstein [17], and the book by Cerf et al. [18].

### 1.1.5 Quantum Networking

A quantum network (see Kimble[19] for a recent review) shares signals among its nodes in the form of quantum information. The building of practical quantum networks is where

---

<sup>4</sup>For the record, this was a somewhat philosophical choice: photon counters were also available. These could also have been used to gain much of the same information using interference measurements. The amplitude information gained by homodyne is just more direct.

much of the experimental effort in quantum information is currently being directed. Apart from communication between the quantum computers of the future, the main reason for this effort is security. The implications of measuring quantum information may be used to guarantee the privacy of a channel.

While such secure quantum communication has been demonstrated (and is indeed already the basis of several commercial devices), a large unsolved problem prevents more widespread use. The problem is distance. The exponential signal decay in even the best transmission medium of optical fibres limits distances to the order of 100km. Techniques have been devised to remove this limitation, using entanglement swapping and a ‘quantum repeater’ which may increase this range. A useful quantum repeater does not yet exist, due largely to the missing requirement of an efficient quantum memory.

Even over short distances where quantum encryption is successful, the technology is mainly used to securely transmit only a classical “key”, which may then be used for highly secure standard communication. This is known as Quantum Key Distribution (QKD).

The process can be summarised briefly, using the common “terminology” of the field. An information sender “Alice” (A) wishes to send information to a receiver “Bob” (B). A third party, “Eve” (E) wishes to eavesdrop on the communication without the knowledge of Alice or Bob. If Alice sends her classical information encoded into quantum variables, any information gained by Eve must necessarily cause a change in the states received by Bob.

## QKD

The problem that is currently faced in (classical) cryptography is not encryption itself; there are many powerful codes for encrypting information, which are considered for practical purposes unbreakable. The difficulty with using these codes is that they require pre-shared information as a starting point, in the form of a password or ‘key’. Such a key requires pre-planning, and is subject to a human factor of being lost or stolen. Further, many modern use-cases require security between two parties who have previously had nothing to do with each other. Perhaps the most obvious example of this is online transactions. Highly successful techniques exist for these situations, but are not proven to be secure (particularly to quantum computers).

A sufficient target of quantum encryption then, is to create a pool of information shared exclusively between two parties which may then be used to encrypt further transmissions. This is known as QKD. See Scarani et al.[20] for a recent review.

There are two main types of protocol for key distribution. The first type is based on the Heisenberg uncertainty principle. The best known example of this is known as the ‘BB84’[13] protocol.

In BB84, Alice generates a random string of bits, and encodes each in a qubit which she sends to Bob. For each bit she randomly chooses a basis from an agreed set (In the common approach for polarisation qubits, she may encode her polarisation either along 0

or 45 degrees). Bob reads his incoming qubits by also picking one of the bases at random. They then communicate over a standard public line (e.g. over the phone). Alice and Bob tell each other the bases they used (not the measurement results). They keep any results where by coincidence they both used the same basis. For these bits, the result that Bob measured should be the same as Alice sent. If they are not, some information has been lost somewhere. This is assumed to have been retrieved by Eve. To check the level of loss, Alice and Bob announce a subset (over a public line) of the qubits they think they've shared. The ratio of discrepancies to good results puts a limit on how much information Eve may have gained.

In the second type of protocol, pairs of entangled qubits are used, one half of which is sent to each of Alice and Bob. The first proposal is now known as the E91[21] protocol. By performing measuring their part of each entangled pair, Alice and Bob gain a bit of shared information each time due to the correlated outcomes. The way in which the entanglement is first shared may vary, but the most direct way is to send photons over a line as for the BB84 protocol. In this case the same limitations are faced. The gain of using entanglement however is that the task can be split into smaller, much more efficient steps using entanglement swapping and "quantum repeaters" (see section 1.1.6).

### Dense coding

A final application of quantum networking may be so-called 'super-dense coding'. By sending one half of entangled qubits instead of bits, a 2-fold increase in the rate of information transfer over a given channel is theoretically possible. That is two classical bits can be retrieved for every qubit sent. At present, it is still much easier to make two classical channels than to employ dense coding. However, the possibility of dense-coding is informative to understand the limits of classical information transmission. After all, it is quantum information (known as 'noise' for classical purposes) which limits the number of clearly distinguishable bits that may be sent down a line.

#### 1.1.6 Quantum Repeaters

The specific problem targeted by a quantum repeater[22] is increasing the rate at which entangled pairs can be created over a long distance. The motivation is to allow quantum-information processes, such as teleportation to be performed over a long distance. Elements of a quantum repeater have been demonstrated[23, 24, 25], but not yet a practical device. One of the biggest limiting factors at present is the performance of present quantum memories. Repeaters are probably the nearest-term motivation for developing good quantum memory.

Directly transmitting photons incurs a loss that scales exponentially with the distance,  $L$ , such that the probability of success decreases as  $10^{-\gamma L}$ . Even for the best optical fibres at telecommunications wavelengths, a loss of 0.2dB/km is incurred. For 100km, this is 20dB, or a success probability of loss of  $10^{-2}$ . This level might be managed by simply

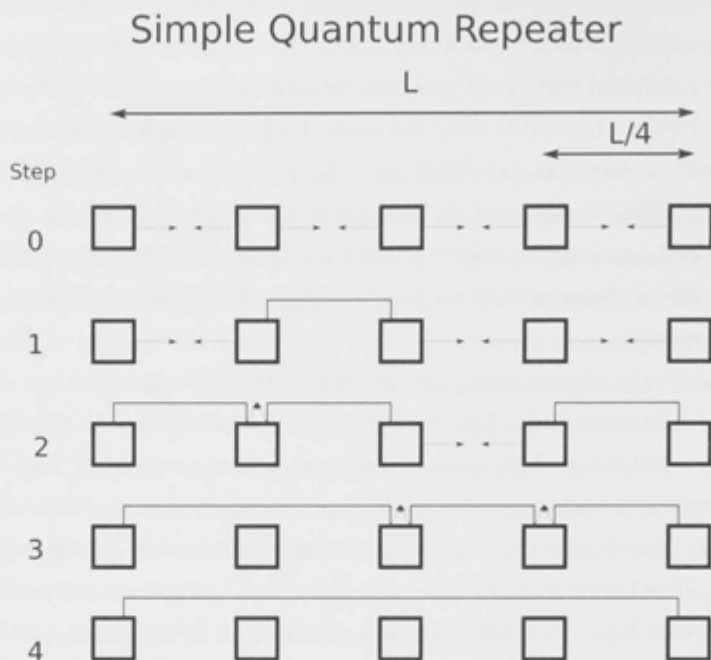


Figure 1.2: A quantum repeater circumvents the exponential scaling loss of fibres by dividing a large length  $L$  into  $N$  sections, with  $N + 1$  nodes between them. Entanglement may then be shared between each node with a much higher success rate. The entangled states are stored in quantum memories. When adjacent links have successfully shared entanglement, they may be joined by a local measurement that cannot distinguish between the links (entanglement swapping). The entanglement is then extended outward step by step. The range of the network may be increased by realising that the entanglement swapping measurement does not have to be performed at the site of each node. Thus each node in this diagram could be split in two and the measurement performed at an intermediate site.

increasing the transmission rate (by a factor of 100). For 1000km however, it's 200dB, or success probability of  $10^{-20}$ . One would need to generate and send entangled pairs at  $R = 10^{20}$ Hz to achieve a transmission rate of 1Hz, even before accounting for other losses.

The principle of a quantum repeater is to divide a large link of length  $L$  into  $N$  sections, with a repeater at each node. Photons carrying entanglement are shared between adjacent nodes with a much higher success rate proportional to  $e^{-\gamma L/N}$ , and stored in quantum memories. Once a node holds entanglement separately with each of its neighbours, it performs entanglement swapping. This consists of a joint measurement of the local information corresponding to each neighbour. If this measurement doesn't distinguish between either neighbour, entanglement is created between the two neighbours. This process is extended in subsequent steps to entangle the whole length. In our diagram (Fig. 1.2 we are imagining 2 memories at each node, but this is not necessary. The entanglement swapping may be performed over a distance, and thus the length may be extended further.

The success of the repeater relies on the ability to ‘herald’ a successful event and to store the appropriate photon until the neighbouring section heralds success. Without heralding, we have a ‘quantum relay’[26] which requires every link to succeed simultaneously. Such a relay does not solve the problem of loss with distance. The form that the heralding takes, and indeed the form that the entanglement transmission takes depends on the protocol details. The details of several practical implementations are given by Sangouard et al. [27]. Some general requirements of these we’ll discuss briefly in chapter 2.

## 1.2 Quantum Optics

As a quantum memory is about preserving the quantum nature of light, we’re concerned with quantum optics. The ‘optics’ we’re concerned with are narrow beams of radiation, i.e. lasers, oscillating at optical frequencies. The field is sometimes considered more broadly to include interactions with atoms, but we save that for the next section.

As this thesis may be read by those who are not quantum-opticians, the purpose of this section is to cover the necessary basics: the properties and types of light that a quantum memory preserves, how they are affected by storage, and how they can be measured.

We overview a ‘standard’ introduction. The free-field is decomposed into *modes*, as expected classically. These modes are then quantised into *states* according to quantum mechanics. Covering these basics allows us to define, in particular, what we mean by ‘modes’. The term is used everywhere in this thesis, so this will hopefully help avoid confusion later, (particularly of the author). Following on, we’ll describe the extreme classes of states- the particle like photon-number states and wave-like coherent states. These form the bases in which more general states may be described, although we don’t discuss this further. Next we discuss the processes of attenuation and amplification which are necessary to understand the limits and effects of a quantum memory (covered in chapter 2). We finally discuss detection, with particular detail on amplitude measurements which we will use in our experiments in chapter 5. The reader is directed to books by Bachor [28] and Walls and Milburn [29] for a proper introduction.

### 1.2.1 Modes of the Electromagnetic Field

In this thesis we consider a slightly artificial concept: modes of the ‘free’ electromagnetic field. We define these as one extreme of a confined field.

It is well known from classical electromagnetism that the field in a reflecting cavity can only form standing waves with certain wavelengths due to self interference. All the possible stable configurations of the field in the cavity can be written as linear combinations of these fundamental modes with various phases and amplitudes.

We consider the limit as the size of the cavity is taken to infinity (in this sense most

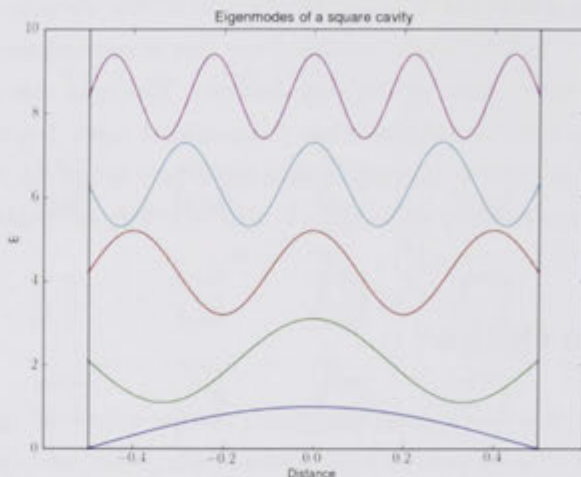


Figure 1.3: When talking about the EM field in free-space, it is common to decompose it into modes of an infinite box. This quantisation scheme is due to Dirac. The choice of basis is arbitrary however, because any mode is allowed by the infinite box.

laboratories have approximately infinite space)<sup>5</sup>. As we do this the spectrum of allowed modes becomes continuous. Combinations of standing wave may now be used to describe an arbitrary shape. When 3 space dimensions are allowed, these fundamental modes are plane waves, and combinations of these can give any observable field: the Gaussian cross-sections of laser beams, spherical radiation from a point source etc. Breaking down free modes in terms of these plane waves is sometimes called Dirac's quantisation.

Plane waves of fixed frequency and unlimited spatial extent are only one choice for a basis. With our infinite basis set, we can construct any well-defined function, and any well defined function may be considered part of a basis function for constructing others. We might reverse our initial thinking and consider plane-waves as a linear combination of non-stationary, travelling pulses of arbitrary shape. Any well-defined envelope in space and time is a candidate for being called a mode. The condition for any 2 such modes to be orthogonal remains the same, that their overlap integral in the space and time is zero. This more general definition of modes was first discussed by Titulaer and Glauber [30], and further elaborated recently by Smith and Raymer [31].

So when we talk about a mode we mean specific spatial and/or temporal profile of light. The modes we usually encounter are those that we generate using a laser. Thus in time we have pulses, and in space we have a Gaussian beam profile. Two laser beams may be 'mode-matched' spatially if they the same beam profile, propagation direction etc. in space. If they are also pulsed synchronously, they are also mode-matched in time.

<sup>5</sup>It is more usual to describe our free modes as those having leaked out of an infinite cavity, rather than still being in one. The principle however is the same.

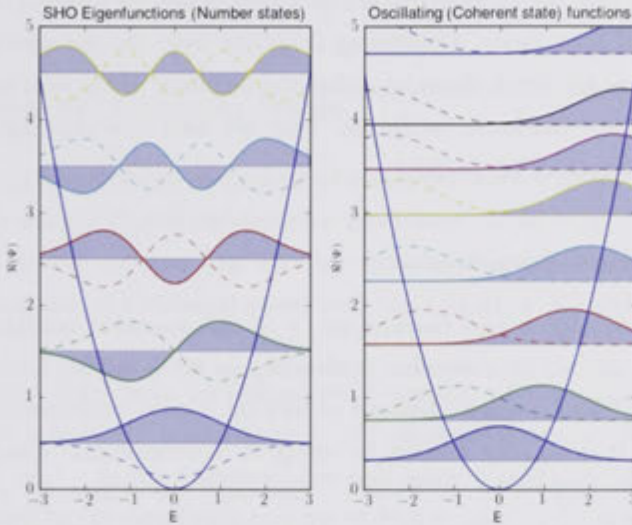


Figure 1.4: A single mode of the field is further quantised according to quantum mechanics into states of the simple harmonic oscillator (in field amplitude, not space). a) Energy eigenstates, corresponding to photon occupation. They occur only at discrete energies. b) Solutions with an oscillating behaviour are also possible, called coherent states. Their energy is not definite. Unlike the eigenstates, their spectrum is not unique and any mean energy is possible. We show here a random selection of mean energies. For both graphs, the dotted lines represent the wave-function half an oscillation (a phase of  $\pi$ ) later.

### 1.2.2 Quantization of a Single Mode

Having defined a mode, we may wish to measure some property of the field in it. When we try to do this with high sensitivity, we run into quantum mechanics. The energy in the field is proportional to the field squared, and is described by a Hamiltonian with the same form as the harmonic-oscillator<sup>6</sup>. Only the ‘motion’ is occurring not in space, but in field amplitude. Thus a single mode of the field is quantised in the same way as the eigenstates of the quantum harmonic oscillator. This is sometimes known as ‘second quantisation’.

The quantum harmonic oscillator is discussed in depth in any introductory textbook (e.g. [32, 33]). The Hamiltonian is written as  $H = a^\dagger a + 1/2^7$ , where  $a$  and  $a^\dagger$  are the ladder operators, referred to as annihilation and creation. They are functions of the position  $x$  and momentum  $p$  as  $a = \sqrt{\frac{m\omega}{2\hbar}}(x + \frac{ip}{m\omega})$ .

The first few eigenstates of  $H$  are shown. These states of definite, equally spaced energies suggest the concept of occupation by light particles— photons. The number operator  $N = a^\dagger a$  is the observable corresponding to the number of photons in the mode.

<sup>6</sup>The ‘standard’, i.e. bosonic, harmonic oscillator that is, photons being bosons. The fermionic harmonic-oscillator may describe fermionic fields

<sup>7</sup>The fermionic oscillator is usually written  $H = c^\dagger c - 1/2$

As with modes, the energy, or photon number, eigenstates are only one class of states. This must be the case as these particle-like states don't correspond at all to what we observe in classical light fields. Oscillating solutions may also be constructed, called coherent states. These are much closer to a classical intuition of the field as a wave, and also what we expect from a harmonic oscillator. They are used as a more appropriate basis for describing wave-like states of the field.

### 1.2.3 Quadrature Amplitudes

We defined the position,  $x$  and momentum,  $p$  of the quantum oscillator. For the field however, these terms can be somewhat misleading as we're talking here about oscillation in field amplitude, not in space. Instead of writing  $a$  as a combination of position and momentum then, it is usually written in terms of dimensionless "quadratures"  $X_1$  and  $X_2$  such that  $a = \frac{X_1 + iX_2}{2}$ ,  $a^\dagger = \frac{X_1 - iX_2}{2}$ . Generally we define  $X_\theta = (e^{i\theta}a^\dagger + e^{-i\theta}a)$ , where  $X_1, X_2$  correspond to  $\theta = 0, \pi/2$ . Our quadrature operators are now on equal footing. They are the operator form of the real and imaginary components of the field, i.e. components that are  $\pi/2$  out of phase, or 'in quadrature'. They obey commutation relations preventing the simultaneous exact measurement of both  $[X_1, X_2] = i\hbar$ . This is analogous to the position-momentum uncertainty relation.

We will be concerned with the uncertainty in measurements of these quadrature amplitudes. As the vacuum state corresponds to the minimum uncertainty  $\langle \Delta X_{1,2} \rangle = i\hbar$ , we typically neglect the factor of  $\hbar$  and use units of vacuum noise. It is also often convenient to use the variance (dispersion), which is then also normalised to  $1^2$  units of vacuum noise.

The quadrature operators also form the basis of 'phase-space' representations of light, in particular the Wigner and Q functions.

### 1.2.4 Photon Number States

The natural eigenstates (like 'standing waves', but in amplitude, not space) of the harmonic oscillator correspond to discrete excitations of the EM field. They are commonly called Fock states or Number states, which are thought of as containing a definite number of photons. The main one which concerns us is the single photon. The single photon forms the basis for many quantum information applications, and is used to demonstrate many quantum concepts. The Fock states generally are defined by  $N|n\rangle = n|n\rangle$  (recall  $N = a^\dagger a$  is known as the number operator). This includes the vacuum state as the special case where  $n = 0$ .

Because of the definite energy, they are the most particle like of states, and thus the most different from the concept of light as a wave. Their energy is completely defined, but their phase is completely random. The amplitude is poorly defined as a result of the phase uncertainty, such that  $\langle \Delta X_\theta \rangle \propto \sqrt{n}$ . The variance scales linearly with the number of photons. This can be seen looking at Fig. 1.4, showing that a number state has no



preference for positive or negative amplitude and that the amplitude probability has no time dependence.

### 1.2.5 Coherent and Gaussian States

The closest quantum mechanical equivalent of a classical wave is a coherent state, which is precisely the type of light that is emitted by a good laser. A coherent state is also emitted by a cold brick, as the vacuum can also be considered a coherent state. A coherent state may be visualised as a classical wave, with the quantum uncertainty of the vacuum superimposed.

Thus a coherent state may be defined as a vacuum with a displaced mean amplitude/phase. The amplitude uncertainty in all coherent states is that of the vacuum  $\langle \Delta X_1^2 \rangle = \langle \Delta X_2^2 \rangle = 1$ . In the (classical) limit of large amplitude, the vacuum noise is negligible compared to the amplitude and we recover a classical description of the field.

Such light is in an eigenstate of the annihilation operator  $a$ , with eigenvalue  $\alpha = \frac{1}{2}(\langle X_1 \rangle + i\langle X_2 \rangle)$ :

$$a|\alpha\rangle = \alpha|\alpha\rangle \quad (1.3)$$

The annihilation operator,  $a$ , is non-hermitian. It has continuum of eigenstates corresponding to the continuum of complex amplitudes that may be had by the field. Strictly speaking, none of these are orthogonal. In practice however, they may be considered as orthogonal if their eigenvalues differ by more than their uncertainties. Thus they can be used as a (over-complete) basis to describe other states, particularly those with wave-like properties. This is the principle of the Wigner, Huisimi ‘Q’, and Glauber ‘P’ descriptions of the field[29]. The concept is used implicitly in e.g. section ??, but is generally not important for this thesis.

Coherent states are often generalised into ‘Gaussian’ states, defined as those on which measurements of any field quadrature be normally distributed. Thus the natural representation for a Gaussian state is a covariance matrix (in addition to the mean):

$$C = \begin{pmatrix} V_1 & V_{12} \\ V_{21} & V_2 \end{pmatrix} \quad (1.4)$$

where  $V_{12}$  represents correlations between the quadrature variations, and is zero for a coherent state. Gaussian states generally are not pure states, incorporating some classical noise. A Gaussian state is a common model for the output of real lasers with some classical noise. This is why we are interested in them.

We note that an oft-encountered class of Gaussian states is the squeezed states. They may maintain the Heisenberg limit  $\langle \Delta X_1^2 \rangle \langle \Delta X_2^2 \rangle = 1$ , like coherent states, but they balance a reduced uncertainty in one quadrature against an increase in the other. They are important in many quantum information applications but not used in this thesis.

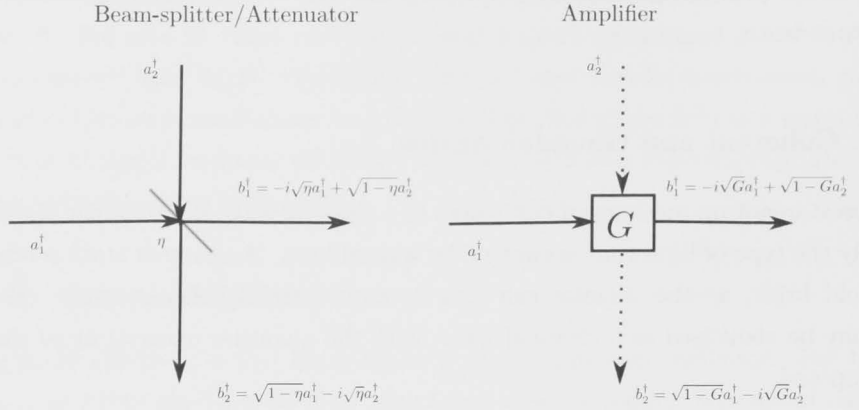


Figure 1.5: A schematic of a beam splitter as a model for attenuation, and a comparable amplifier acting at the quantum limit.

### 1.2.6 Beam-splitters and Amplifiers

Many things in quantum optics are understood as the mixing of a signal mode with one or many auxiliary modes,. Each input mode is assigned its own annihilation and creation operators  $[(a_1, a_1^\dagger), (a_2, a_2^\dagger)\dots]$ . The output operators  $[(b_1, b_1^\dagger), (b_2, b_2^\dagger)\dots]$  are written as combinations of those of the input to fully describe the mode mixing. The auxiliary input modes are commonly ‘empty’, that is occupied only by the vacuum state, but in principle any state is possible.

Attenuation and amplification are both examples this. An attenuator is represented by a beam splitter combining a signal mode with the vacuum, and the amplifier by a similar (although more generic) device.

For an ideal beam-splitter with transmissivity  $\eta$ , the transmitted mode becomes (ignoring the phase factors)[34]:

$$\mathbf{a}_{out} = \sqrt{\eta}\mathbf{a}_{in} + \sqrt{1-\eta}\mathbf{a}_{aux}. \quad (1.5)$$

For ideal linear amplification with gain  $G$ , the output mode is (again ignoring any phase):

$$\mathbf{a}_{out} = \sqrt{G}\mathbf{a}_{in} + \sqrt{G-1}\mathbf{a}_{aux}^\dagger. \quad (1.6)$$

For the amplifier, the physical origin of the auxiliary ports depends strongly on its implementation, but is more often corresponds to the inner state of the amplifier rather than windows for extra light to enter and exit. The stronger the amplification, the more of the auxiliary mode is mixed in.

If the auxiliary ports are empty, the quadrature variances of the outputs become, in

units of vacuum state variance

$$V_{1,2}^{out} = \eta V_{1,2}^{in} + (1 - \eta) \quad (1.7)$$

for linear attenuation, and

$$V_{1,2}^{out} = G V_{1,2}^{in} + (G - 1) \quad (1.8)$$

for amplification.

The beam splitter is one of the most important devices in quantum optics, not only describing attenuation, but providing a model system any time modes are split. A quantum limited amplifier is another useful conceptual device, and is closely related to measurement. It may be used as an aid in optimal copying or measuring of quantum states. We'll look at this in chapter 2.

### 1.2.7 Counting Photons

Most optical detectors count photons in some sense, giving an output scaling with the number detected. Incident photons are converted into electrons which appear as a voltage on a recording device. They may be broadly divided according to purpose: the most highly advanced sacrifice some efficiency for low noise and high sensitivity for counting individual photon events. These are known as Single Photon Detector (SPD)s. On the other hand, more robust intensity detectors have a higher noise floor but also higher quantum efficiency. Their purpose is to measure photon flux. These form the basis for amplitude measurements as we'll describe next. The description 'photon-counter' is more often applied only to SPDs (although until recently, many 'counters' only 'counted' around 1 photon in 10).

As the vacuum state has a definite photon number of zero, a photon detector may potentially have a very low background noise level. This is true at least at optical energies where the thermal background is negligible, although the same cannot be said at RF frequencies.

A SPD is the natural method for detecting Fock states. The detector cannot distinguish phase but the Fock state has no definite phase to detect. Phase-sensitivity can be introduced by combining the state to be measured with a reference state, and interference fringes measured. In the extreme of using a very large coherent state as a reference, an amplitude detector may be realised. This is known as homodyning (also heterodyning) and is the topic of section 1.2.8.

### Splitting the mode into time bins

Usually a detector is run, and a signal is recorded, continuously in time. This may be thought of this as taking many measurements of many short-time modes (as narrow as the detector bandwidth allows). But we're generally interested in a longer time window-

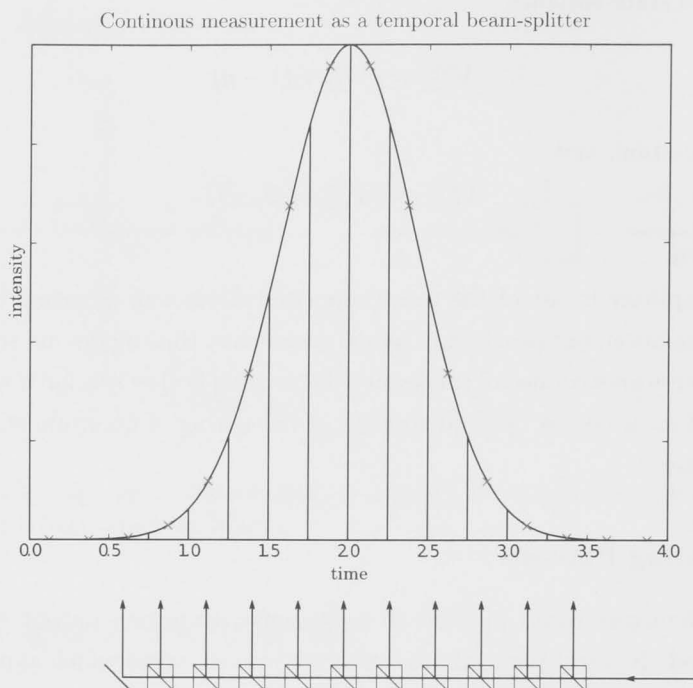


Figure 1.6: An intensity vs time trace of a mode to be measured can be thought of as temporal beam-splitting, dividing the pulse into many temporal modes. As we are measuring energy, which is a conserved quantity, the sum of all signal in all the temporal modes is the same no matter how it is divided.

perhaps a longer pulse, or an essentially infinite time if trying to measure a frequency side-band.

Of course classically there is no issue- if we want to know how much energy is in a pulse or side-band, we simply sum over the pulse time, or the frequency range of interest. We can also do the same in the quantum case, because energy and photon number are conserved even when a signal is divided. An intuitive way to understand the measurement is to compare it to a system of (temporal) beam splitters, each with an empty auxiliary input port. If we measure all the outputs and add up, we get the same total number (or the same total energy) regardless of how many times we divided it.

### 1.2.8 Measuring Amplitude

Homodyne and heterodyne detection are the main tools for measuring quadrature amplitudes. These are described in standard quantum optics textbooks, e.g. Walls and Milburn[29]. Our main aim is to describe the aspects of the technique that we'll need for experiments in chapter 5, and for a discussion of 'classical' memory in chapter 2. We give some detail on the practical aspects of time-domain techniques as opposed to the more common frequency domain. The final novel aspect of the discussion is that we consider

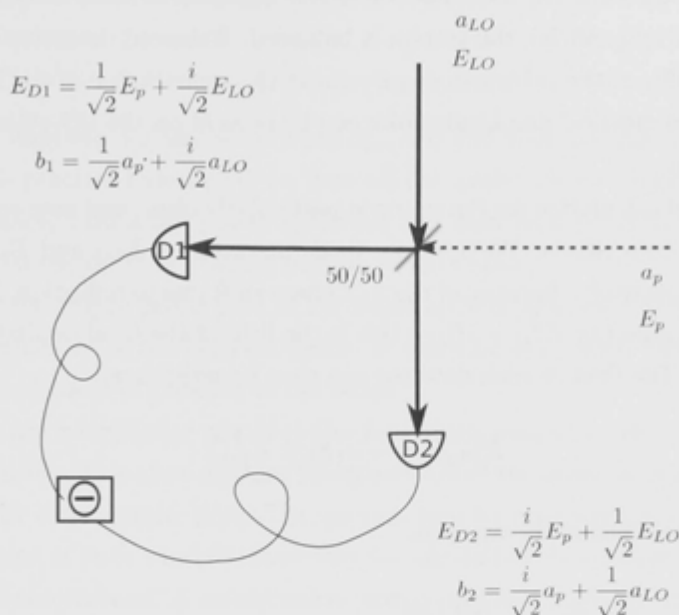


Figure 1.7: Balanced homodyne/heterodyne detection.

doing much of the analysis in post-processing. This reduces the requirements on detectors.

### Homodyne and Heterodyne

The amplitude of an optical field is also measured using intensity detectors, but indirectly through the techniques of homodyne and heterodyne.

There are two big problems facing a direct measurement. The first is that optical fields oscillate at  $\approx 10^{15}$  Hz, much faster than any detection electronics can go. The second is that the quantum fluctuations we are interested in are generally much smaller than common detection noise sources.

These problems are solved by mixing the field to be measured (the probe) with a much stronger reference beam called the Local Oscillator (LO). When beat against the large field of the LO, the small amplitude variations become large intensity variations in the combined beam. This signal will oscillate only at the difference frequency, which may easily be very small or zero. This is especially easy if the probe and LO are derived from the same source.

The sensitivity scales linearly with the amplitude of the LO, and this field is (ideally) large enough that it may be treated classically. This allows us to think of the LO as an exceptionally strong linear filter. Our detection is only sensitive to fields in the same mode and exactly  $\pi/2$  out of phase with it—any signal which doesn't add coherently with the LO is too small to register.

To realise the mode-matching in practise, the beams are combined at a 50/50 beam

splitter. Both outputs are measured, and the signal from each is subtracted. If the beam-splitter is truly 50/50, the system is balanced. Balanced detection has a significant additional benefit, as the subtraction can remove the contribution of the LO's intensity to our signal. Thus classical amplitude noise can be present on the LO without effecting the measured signal.

The classical calculation for the set-up is particularly clear, and may easily be extended to include quantum effects. We'll let our fields be given by  $E_{LO}$  and  $E_p$  for the LO and probe fields respectively. Because of the  $\pi/2$  phase shift due to reflection, it can be simpler to consider the quantity  $E_{LO}^{\frac{\pi}{2}} = iE_{LO}$ ; this is the field of the local oscillator phase shifted by 90 degrees. The field at each detector can then be written as

$$E_{D1,D2} = \frac{1}{\sqrt{2}}(E_p \pm E_{LO}^{\frac{\pi}{2}}) \quad (1.9)$$

The measured intensities are then:

$$I_{D1,D2} = |E_{D1,D2}|^2 \quad (1.10)$$

$$= \frac{1}{2}(|E_{LO}|^2 + |E_p|^2 \pm (E_p E_{LO}^{\frac{\pi}{2}*} + E_p^* E_{LO}^{\frac{\pi}{2}})) \quad (1.11)$$

If the beam-splitter is exactly balanced, and the detectors are perfectly matched, when we subtract the signal from our two detectors the first terms cancel while the cross-terms survive. The signal we measure is the difference in intensity between the two detectors:

$$I_{sub} = (E_p E_{LO}^{\frac{\pi}{2}*} + E_p^* E_{LO}^{\frac{\pi}{2}}) = 2\Re E_p E_{LO}^{\frac{\pi}{2}*} = 2|E_p||E_{LO}^{\frac{\pi}{2}}| \cos(\phi) \quad (1.12)$$

$$= 2|E_p||E_{LO}| \sin(\phi) \quad (1.13)$$

were  $\phi$  is the phase difference between the beams. Generally we require that  $E_{LO} \gg E_p$ . As we can see, the (classical) intensity of the LO does not appear. Variations in  $E_{LO}$  may still have an effect, but only if they are large as a fraction of the full field strength. Thus we have a detector which measures the component of  $E_p$  in the phase defined by  $E_{LO}$ .

The distinction between homodyne and heterodyne is simply whether the beams are exactly matched in frequency or not. Homodyne corresponds to the case where the signal field has the same frequency as the LO so that they maintain relative phase. In practise the phase can be chosen by varying the path length of the LO, e.g. using a Piezo mounted mirror or similar. In a heterodyne measurement, the frequency difference causes the relative phase to scan at the beat frequency,  $\phi(t) = \Delta\omega t + \phi_0$ . A heterodyne measurement is thus alternating between measuring the in-phase and in-quadrature components of the signal.

Homodyne makes the best measurement of a single quadrature, while heterodyne gives us a measure of both. As there are two of them, each heterodyne sample cannot be as accurate the single homodyne one. The differences are discussed very briefly in section 1.2.8. We now look at how the quantum noise appears in measurements, and then how we

extract a measurement of the quantum amplitude in the single mode defined by a pulse.

### Shot Noise

In our classical analysis, we can simply put  $E_p = 0$  and then no signal will be seen on our detector. In practise however, if we turn off our probe we see quantum noise. This noise is white noise, with a flat power spectrum and with each point in time completely uncorrelated from the next. It's called vacuum or shot-noise.

If we ignore the empty probe port of the beam splitter, the noise may be interpreted as intensity fluctuations of the local oscillator. In this case it is called "shot-noise" due to the random arrival times of photons in the beam.

However we are interpreting our device as one which measures the field in the probe, not the LO intensity. This gives another interpretation of the noise, as due to the vacuum state entering the empty probe port. The vacuum field we may write  $E_p(t) = \delta_0(t)$ .

The properties of such random noise can be unintuitive but follow simply from its complete lack of correlations. A measurement at one time is completely uncorrelated from a measurement at any other.

To explore the properties, let us imagine a digital oscilloscope connected to a perfect homodyne detector that is measuring the vacuum. This displays a time trace of the intensity that is incident. As the oscilloscope is digital, the trace is discretized into time bins separated by  $\Delta t$ . Each point corresponds to a measurement of the average intensity in that interval. The points deviate slightly from zero with standard deviation  $\delta_{\Delta t}$ . This is most easily calculated by considering it as shot-noise of the local oscillator. Each photon striking the detector is a completely random event so that the arrival statistics are poissonian. In this case fractional deviation for an average of  $\langle N \rangle$  arrivals is simply  $1/\sqrt{\langle N \rangle}$ . The expected deviation of points on the oscilloscope is then:

$$\delta_{\Delta t} = I_{LO}/\sqrt{\langle N \rangle} = I_{LO}/\sqrt{\frac{\Delta t I_{LO}}{\hbar\omega}} \quad (1.14)$$

$$= \sqrt{\frac{I_{LO}\hbar\omega}{\Delta t}} \quad (1.15)$$

Interpreted as a measurement of the vacuum, the signal we see is due to  $2E_{LO}E_{vac}$ . Thus the apparent field and intensity of the vacuum can be calculated by attributing the fluctuations to  $E_{vac}$ :

$$2\delta(E_{LO}E_{vac}) = \delta_{\Delta t} \quad (1.16)$$

$$\delta E_{vac} = \frac{1}{2}\sqrt{\frac{\hbar\omega}{\Delta t}} \quad (1.17)$$

$$(1.18)$$

and also

$$I_{vac} = \frac{\hbar\omega}{4\Delta t}$$

This has some curious properties. The intensity of the vacuum that we measure depends on  $1/\Delta t$ , that is on the measurement rate. The faster we measure, the higher the apparent intensity. In different terms, the intensity measured is proportional to the bandwidth of the measurement. This is the cause of the white spectrum for shot noise.

On the other hand, the mean *energy* in any given time interval is  $I_{vac}\Delta t = \frac{\hbar\omega}{4}$ . That is, no matter what time we sum over, the mean energy measured in the vacuum is  $1/4$  that of a photon.

The signal averaged over any given interval can be considered measurement of the vacuum in the mode represented by the time window  $\Delta t$  (and the spatial mode and phase determined by the local oscillator beam). Because of its consistency, it is convenient to normalise everything in an experiment to the shot noise, such that  $I_{vac} = 1$ . Another convention is to the energy of the photon, such that  $I_{vac} = 1/4$ . In this thesis we'll usually follow the former.

With noise included, the measured signal is  $E_p(t) = \langle E_p(t) \rangle + \delta(t)$  with  $\delta(t)$  potentially being any kind of probability distribution. In general,  $\delta(t)$  will include both noise due to measurement and that due to the state. The part that is due to the state, together with the mean, completely defines the state.

### Pulse measurements

Often our signals come in the form of a pulse. Each time a pulse is recorded, we wish to obtain a single value of the field *quadrature* of the pulse, that is a measurement of a quadrature operator. Classically the problem is equivalent to applying an optimal filter to a noisy signal, to emphasise the parts which contain the most information. The optimal filter is the mode envelope, which we write  $S(t)$ . We assume it is normalised such that  $\int_{-\infty}^{\infty} S(t)^2 dt = 1$ . In terms of  $S(t)$ , the signal pulse is  $\langle E_p(t) \rangle = (E_p^{rms})S(t)$ . The energy in the pulse is  $(E_p^{rms})^2$ .

The common way to make this measurement is to pulse the LO in the same mode as the signal, thus using the LO as the filter in time (as it already is in frequency and in space). The area in the signal is then integrated to give a single number. The local oscillator field may be written  $\langle E_{LO}(t) \rangle = (E_{LO}^{rms})S(t)$ .

The signal intensity we expect at the output of an ideal homodyne detector is

$$\langle I_{sig} \rangle = 2S(t)^2 E_{LO}^{rms} E_p^{rms,\phi} \quad (1.19)$$

where we now only consider the component of  $E_p$  in phase with the LO.

The (un-normalised) quadrature amplitude in the mode defined by  $S(t)$  is given by the area of the signal:

$$\langle X_\phi \rangle = 2E_{LO}^{rms} E_p^{rms,\phi} \int_{-\infty}^{\infty} S(t)^2 dt = 2E_{LO}^{rms} E_p^{rms,\phi} \quad (1.20)$$

$$(1.21)$$



Thus we simply need to know the energy of the local oscillator pulse to get our normalised value.

As a check we calculate the uncertainty we would expect in such a measurement of a coherent (or vacuum) state. This is simple as the noise at different times is uncorrelated, so that the variances sum. Let  $\sigma_0 = \langle \Delta \delta_0 \rangle$ , the noise (standard deviation) of the vacuum. This is again filtered by the mode envelope, and integrated over time. The fact that the variances add simplifies the integral, giving:

$$\langle \Delta X_\phi \rangle^2 = \int_{-\infty}^{\infty} (2E_{LO}^{rms} S(t) \sigma_0)^2 dt = (E_{LO}^{rms})^2 \sigma_0^2 \quad (1.22)$$

$$\langle \Delta X_\phi \rangle^2 = 4E_{LO}^{rms} \sigma_0 \quad (1.23)$$

The signal-to-noise ratio when estimating the amplitude of a coherent state is then  $\langle X_\phi \rangle / \langle \Delta X_\phi \rangle = E_p^{rms, \phi} / \sigma_0$ , independent of the mode shape. Thus the measured noise of the vacuum may be used to normalise a measurement even if LO intensity and the detector gains aren't known.

Many high-performance homodyne systems use a pulsed LO like we've described, combining it with fast analog electronics to rapidly process the signals. This way measurements can be made very quickly, up to many million per second.

It should be clear however that the reason for pulsing the LO is speed and convenience. Its role is essentially a temporal filter to amplify the parts of the signal we're interested in. So long as our detection system is capable of clearly recording the pulse, we can achieve the same measurement by using a constant LO and performing the filtering in post-processing. This is slower but potentially simpler and neglects the need to know the envelope at the time of measurement. Further, there is a slight gain in that performing the filtering after the measurement also filters out other noise sources such as electronic noise of the detectors.

As a comparison, one could potentially do the same for spatial mode-matching, using a much broader beam diameter for the LO than the signal. To do this however, we would require spatial sensitivity in our detection, such as a CCD detector. While this is probably not worth the effort, the opposite is sometimes used: using a small LO (small in space) to pick out parts of, and thus characterise, the spatial mode of the probe[35].

## Digital Balancing

Implementing well balanced homodyne detection is a significant technical challenge. Lasers typically have increasing classical noise levels towards low frequency, precisely where most of a homodyne signal usually is<sup>8</sup>. To get around the noise, suppression of intensity noise better than 40dB is often achieved[36]. This requires that the signal level measured on

---

<sup>8</sup>Sometimes homodyne is performed on side-bands centred around, but not at, zero detuning. In this case low frequency noise is less of an issue

each detector is matched to better than 1 part in 10000. In order to achieve this precision, the first challenge is in designing and building high quality, perfectly matched detectors. The second is matching the signal seen by each.

The point we make here is that if well matched detectors and good 50-50 beam splitters aren't available, good measurements are still possible. As long as we can measure the signal at each detector accurately, we may account for these imperfections in post-processing. The main limitation becomes the precision and dynamic range of the measurement. A further caveat of not having the same light level on both detectors is a stricter requirement that  $|E_p| \ll |E_{LO}|$ .

After recording a (unbalanced) signal, we simply multiply the output from each detector by a different gain factor  $G_{1,2} = g_{1,2}(\omega)B_{1,2}$ . The gain factors  $g_{1,2}(\omega)$  are complex functions that account for the frequency dependent gain of each detector. These may be obtained simply by measuring the response of each detector relative to a known good source, or to shot-noise. The second factor ( $B_{1,2}$ ) accounts for an unbalanced beam splitter, such that a different light level is incident on each detector. If the beam-splitter has reflectivity  $\eta$ , we simply set

$$B_1 = \frac{\sqrt{\eta}}{\sqrt{1-\eta}}, B_2 = \frac{1}{B_1}. \quad (1.24)$$

We may confirm that our signal remains the same as the unbalanced case (so that we're not sacrificing our signal to achieve good balance). We start with the intensity measured at each detector in these conditions:

$$I_{D1} = (1 - \eta)|E_{LO}|^2 + \eta|E_p|^2 - 2\sqrt{\eta}\sqrt{1-\eta}\Re(E_{LO}E_p^*) \quad (1.25)$$

and similar for  $D_2$ . The subtracted signal is then  $I_{sig} = B_1I_{D1} - B_2I_{D2}$ , where we have assumed that the gain factors  $g_{1,2}(\omega)$  (if necessary) have already been applied. The measured signal is

$$I_{meas} = |E_{LO}|^2(B_1(1 - \eta) - B_2\eta) - 2(B_1 + B_2)(\sqrt{\eta}\sqrt{1-\eta})\Re(E_{LO}E_p^*) \quad (1.26)$$

We have neglected the  $|E_p|^2$  which would normally be cancelled with a 50/50 beam splitter, assuming that it is much smaller than  $E_pE_{LO}$ .

We want the LO classical noise to be balanced so that the first term in (1.26) disappears. Thus we require

$$B_1(1 - \eta) - B_2\eta = 0 \quad (1.27)$$

The second term in (1.26) represents the measured signal. We want this to be independent of  $\eta$ , which it will be if

$$(B_1 + B_2)\sqrt{\eta}\sqrt{1-\eta} = 1 \quad (1.28)$$

Both these requirements are satisfied by substituting the given values for  $B_{1,2}$ .

Given that the signal we measure remains the same, so must the measured noise level.

This can be seen by considering the intensity ‘shot’-noise due to the LO on each detector. The noise on each detector is independent and proportional to the intensity, giving us

$$B_1^2(1 - \eta) + B_2^2\eta = 1 \quad (1.29)$$

This is also satisfied by our given  $B$ s.

In practise we want a ‘knob to twiddle’, or a parameter for a computer to optimise, in order to fine tune the balance. The observed noise level provides a parameter to minimise. We maintain  $B_1B_2 = 1$ , while varying  $B_1$  to minimise observed noise (or a classical signal) on the local oscillator. Under these conditions, both the quantum (1.29) and classical noise contributions to our signal (1.27) are minimised at the correct values. (After the minimisation we could use the values of  $B_{1,2}$  as an accurate measurement of  $\eta$ ).

### Homodyne vs Heterodyne

We have mainly discussed homodyne as this is what we use for quantum noise measurements in chapter 5, and it is the technique more widely used in quantum optics. It may often be easier to perform heterodyne measurements however, in particular because it may avoid the low frequency noise seen on most lasers.

Ignoring practical difficulties, the main difference for quantum measurements is that homodyne gives a good sample of one quadrature, while heterodyne samples both quadratures less accurately. The differences are illustrated in Fig. 1.8, where we compare two orthogonal homodyne measurements, one in each quadrature, to a single heterodyne measurement.

We may consider the extra noise source in the heterodyne measurements in two ways (three if we say “it must be this way to satisfy the uncertainty principle”). In frequency space we see that our detection makes no distinction between positive and negative frequencies, as a direct consequence of measuring only real signals. This means that attempting to measure signals which aren’t symmetric about 0 detuning, we mix in extra vacuum noise. This is illustrated in Fig. 1.8b. Considering the modes in time, we can see that we’re only measuring each quadrature half the time. Thus it’s not that we’re measuring more noise, but less signal.

Homodyne is the more often used technique for quantum measurements. In terms of the number of measurements that must be made, it is the better choice for determining a specific parameter of a state[37]. Some of the information from a heterodyne measurement will simply be wasted.

The techniques are on a more even footing for general tomography. Here an unbiased reconstruction of the entire state is required so that any and all the information obtained by a measurement is of use. From a practical point of view, heterodyne is probably easier too, and not just to avoid low frequency noise. Full tomography using homodyne requires taking many measurements at a distribution of phases (each measurement is a sample of

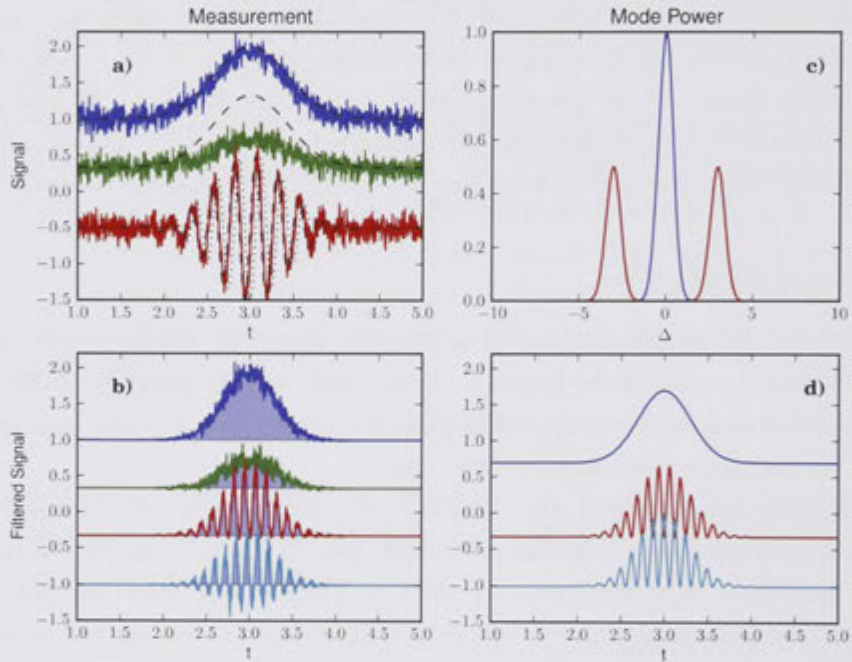


Figure 1.8: Homodyne and heterodyne pulsed-quadrature measurements. a) Simulated measurement traces for weak coherent pulses near the quantum limit. Blue and green lines correspond to separate homodyne detection in orthogonal quadratures. The broken lines are envelopes corresponding to the pulse mode, which we use as filters. For the single heterodyne measurement, there are two envelopes, one for each quadrature. b) (bottom left) The signals after filtering with the mode envelopes. A single number is obtained from each trace by integrating the area. c) Spectral power density. The heterodyne mode is split between positive and negative frequency, and thus includes noise from both, while the signal to be measured is only on one side. d) The mode-filtered intensities in time. This shows that each heterodyne mode only includes half the signal energy of the homodyne case. With this in mind, we might think of each heterodyne sample as having less signal than homodyne, rather than more noise. However, adding the signal energy from both heterodyne samples gives the same as the homodyne mode.

a marginal distribution, i.e. a projection, of the state's Wigner function corresponding to the phase of the measurement[35]). In a truly unbiased measurement, the phases of the measurements should be randomised. A heterodyne measurement has a simple correspondence to the state's Q function, each pair of heterodyne measurement giving a single sample. Thus changing the phase of the measurement is not required.

In any case, heterodyne remains the perfect tool for high-sensitivity classical measurements, where the quantum noise is negligible. It's for this purpose we'll use it in chapter 5. For quantum noise measurements, we will use homodyne.

### 1.3 Atom-light Interactions

A practical way to manipulate light is to first interact it with atoms, then manipulate the atoms. This is how the quantum memories we study are implemented. In this section we give a brief summary of the concepts and results we'll use to understand these memories.

The general problem of electromagnetic radiation interacting with atoms is of course exceedingly complex. A single atom, with all its electrons and internal degrees of freedom is a challenge to understand in itself. That is even before it's put into a crystal environment and interacted with a quantised field varying rapidly in space and time. Fortunately for us, we are only interested in a tiny corner of this complex problem. We study visible light interacting with well isolated transitions, with no quantum correlations between atom and field to worry about.

We begin in 1.3.1 by describing this highly simplified picture. This will serve to give a feel for the problem, as well as to collect in one place the assumptions we will make.

An assumption that we usually, but not always, make is that we have far more atoms than photons, so that atoms remain mostly in the ground state. In this case the interaction is linear and the problem greatly simplified. We may treat the atoms as classical dipole oscillators and radiators. These dipoles are driven by the field, and then re-radiate in proportion to the driving strength. This simple linear-response picture is adequate for a majority of phenomena we will study.

For more intense light fields, we need a simple extension to this model in the form of the semi-classical Optical-Bloch Equations (OBE). The states and evolution of the atoms are conveniently expressed on the Bloch-sphere. We'll use the Bloch-sphere to describe the evolution of an ensemble of atoms, each resonant at a slightly different frequency. These 'inhomogeneously broadened' ensembles are the second main concept required for understanding memories. We could also describe them in the dipole picture.

A more general (and rigorous) description of light propagation is given by Maxwell's wave equation coupled with the atomic polarisation given by the OBE. These contain our linear approximation in one extreme. The resultant Maxwell-Bloch Equations (MBE)s may be further approximated to give a simple form. This is particularly suited to numerical simulations, which we use throughout this work.

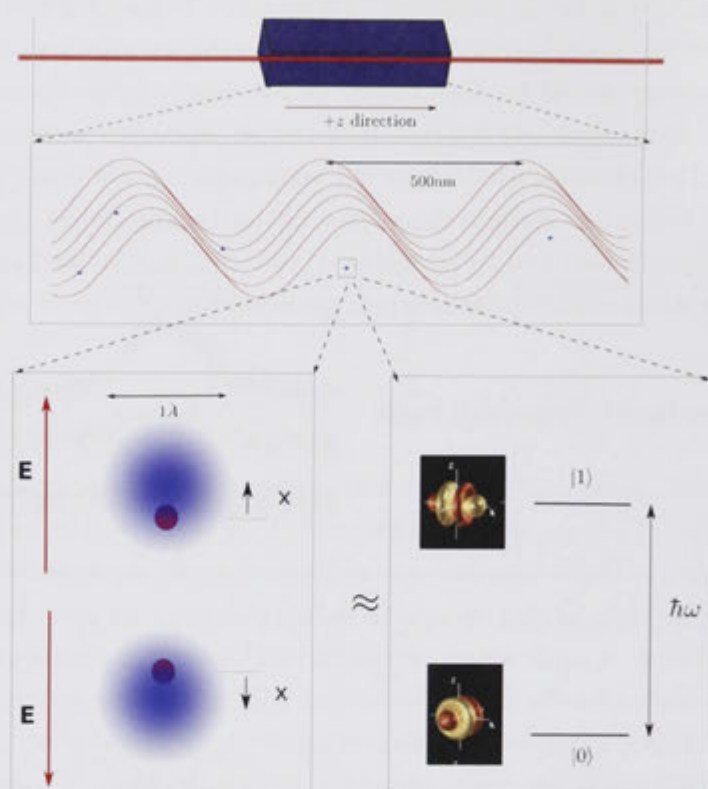


Figure 1.9: The spatial scales for ions in a solid interacting with visible light. The medium is much larger than the wavelength of light, the ions are much smaller. This leads us use the plane wave and the dipole approximations (see text). The atoms themselves we can often treat as classical dipoles, which are a small-excitation limit of a simple two-level model.

We finally discuss phenomena of light propagation. Absorption, dispersion, Free-induction Decay (FID)s and slow light are described by standard time-invariant linear systems theory. Memories ultimately are time-dependent, for which the impulse response can be more helpful than the spectral response considered so far. Simple photon-echo techniques that form the basis of quantum memories are described naturally in these terms.

### 1.3.1 Approximate Picture for Light-Ensemble Interaction

The light we consider is coherent and visible (its a laser beam), with a wavelength  $\sim 500\text{nm}$ . This is 3 orders of magnitude larger than the size of an atom at  $1\text{\AA}$ . This allows the *dipole* approximation to be made. This means we worry only about the time-variation of the field seen by the atom. The gross structure of the atom's electron cloud is represented by its dipole moment. A third spatial scale is given by the diameter of the laser beam. At  $> 100\mu\text{m}$ , it is much greater again than the wavelength. This lets us to consider the laser

as made up of 1D plane-waves.

The optical transitions of the rare-earth ions which we are primarily concerned with in this thesis are exceptionally isolated, interacting only very weakly with their immediate environment. Considering them as oscillators, they have tremendously high quality factors on the order of  $10^{11}$ . At this Q, the 1m pendulum of a grandfather clock would oscillate for around 5000 years. This high Q-factor is responsible for several approximations. Firstly it means that the response drops off very quickly away from resonance, so we only consider light near-resonance. Similarly, we may make the Rotating Wave Approximation (RWA). This effectively averages out the rapid oscillation of our driving field<sup>9</sup>. Finally the spectral narrowness of these transitions (and the narrowness of our laser source) means that only two levels are coupled at one time, so that we can use a simple 2-level model.

The laser light oscillates every  $10^{-15}$ s, whereas we are interested in pulses and time-scales of  $\mu$ s. This allows us to make the Slowly Varying Envelope Approximation (SVEA) for time. Further the coupling of our atoms to the field is small enough that the field is affected negligibly after traversing a distance of a wavelength. This lets us make the SVEA also for space. Mathematically, this involves separating out the slow-varying envelopes from the rapid oscillations as  $E_{in} = E^{env} e^{i\omega t - kz}$  while assuming that  $\frac{1}{E^{env}} \frac{\partial E^{env}}{\partial t} \ll \omega$  and  $\frac{1}{E^{env}} \frac{\partial E^{env}}{\partial z} \ll k$ . Similar results apply to the atomic polarisation.

Finally, there are two regimes of light intensity used in this thesis: when driving atoms, intense light is used which may be considered purely classically as quantum effects are negligible. When considering light for storage in a memory, very weak beams are used. We may again treat the field classically as we are interested only in how its mode transforms. This is because the interaction is linear in this low-intensity regime. The atoms and light, being independently prepared, have no quantum correlations between them.

### 1.3.2 Damped, Driven, Dipole Oscillator

The picture of an atom as a dipole oscillator, also known as the Lorentz model, is probably the simplest model of the atom with any basis in reality- and precisely for this reason, it is very useful. We'll first consider the generic oscillator, then make some observations about its back-action on the field in section 1.3.3

A convenient picture for our dipole is an electron cloud attached to a heavy nucleus by a spring, experiencing a restoring force according to Hooke's law,  $-kx$ . It is velocity damped by  $-b\dot{x}$  and driven by a field representing a force  $F(t)$ . The generic oscillator is

$$\ddot{m}x + b\dot{x} + kx = F(t) \quad (1.30)$$

In this form,  $\gamma = 2b/m$  gives damping, and  $w_0 = (k/m)^{1/2}$  is the resonant frequency

<sup>9</sup>It is equivalent to approximating both the linearly polarised laser and transition as circularly polarised. See e.g. Suter[38]

( $k$  and  $m$  are spring stiffness and mass as usual).  $F(t) = qE(t)$  is an external force, due to the light with electric field amplitude  $E(t)$  influencing the charge  $q$ . We assume that it can be described primarily as an oscillation at  $\omega_0$ , modulated by a slowly varying envelope such that  $F(t) = qE_0e^{i\omega_0 t}$ .

Under the influence of a steady driving field, ( $E_0e^{i\omega t}$ ,  $E_0 = \text{constant}$ ) we assume that the dipole responds at the same frequency with amplitude given by  $x = x_0e^{i\omega t} = \mathcal{X}E$ . Substituting for  $F$  and  $x$  in the equation of motion, the oscillatory factors cancel, and we find that both the phase and magnitude of the dipole's response vary with the driving frequency. The response is given by

$$\mathcal{X}(\omega) = q/m \frac{1}{\omega_0^2 - \omega^2 + 2i\gamma\omega}. \quad (1.31)$$

We are interested only in frequencies close to resonance. We define the detuning of the field from the dipole resonance as  $\delta = \omega - \omega_0$ , and note that in our case  $\omega \sim 10^{15}$ ,  $\delta \ll 10^7$ . It is then a very good approximation to write  $\omega_0^2 - \omega^2 = 2\Delta\omega_0$ . Making this substitution emphasises the form of the above:

$$\mathcal{X}(\omega) = \frac{q}{2\omega_0 m} \frac{1}{\Delta - i\gamma} \quad (1.32)$$

We see that when driven far from resonance, the dipole follows the field ( $\phi = \arg\mathcal{X} = 0$  below resonance,  $\phi = \pi$  above), while close to resonance it lags behind  $\phi = \pi/2$ <sup>10</sup>.

It is common to consider the components in-phase and in-quadrature with the driving field separately (corresponding to resonant and off-resonant driving). We let  $x_0 = u - iv$ , with  $u, v$  representing the slowly varying quadratures of the envelope. We substitute this into the equation of motion, equate real and imaginary parts, then make the slowly-varying envelope approximations of  $\dot{u} \ll \omega u$ ,  $\ddot{u} \ll \omega^2 u$  (and the same for  $v$ ). The result is

$$\begin{aligned} \dot{u} &= -\Delta v - \frac{\gamma}{2}u - g\Im(E_0) \\ \dot{v} &= \Delta u - \frac{\gamma}{2}v - g\Re(E_0) \end{aligned} \quad (1.33)$$

where  $g$  ( $= q/\sqrt{2m\omega}$  here) is some indicator of the dipole-light coupling strength. We're merely interested in the simple picture here, and don't expect to calculate  $g$  from it.

In general there is not a single atom (or oscillator), but an ensemble. The field interacts with them simultaneously, so that they appear as a single, strongly coupled oscillator. We are interested in the polarisation density,  $P = Nx$  where  $N$  is the atomic number density. The polarisation response to the field is then given by  $\chi = N\mathcal{X}$ .

These equations are the small-pulse limit of the semi-classical OBEs (see 1.3.4). This small-pulse limit corresponds to many situations of interest.

<sup>10</sup>As may be grasped intuitively by playing with a mass attached to a spring.



### 1.3.3 Re-radiation by Damped Dipoles

The above is assuming only one source of electric field: a laser. The dipoles will also create a radiation field of their own which will add coherently with the laser. This back-action on the field becomes significant as the number of oscillators, or the interaction strength of each, increases. It leads to absorption, dispersion and the other phenomena of light-matter interaction.

The problem of how strongly interacting dipoles radiate is complicated, because in general we can't separate which is the external field and which is that due to the atoms. But we can get the form of how a field is affected by the atoms without knowing the details — we just need to know that the response is linear in amplitude. We'll work out the exact size and phase in 1.3.6.

Let us assume that when a CW field is applied, the output is linearly related in (complex) amplitude, but at the same frequency. The output that would be expected with no medium present we'll call  $E_0$ . Any deviation from  $E_0$  is ascribed to coherent radiation by the medium. In terms of slowly varying envelopes this may be expressed as

$$E_{out} = E_0 + E_r = E_0(1 + H) \quad (1.34)$$

$H$  is the spectral response of the medium as a whole. Negative  $H$  means that radiation of the medium cancelled some of the incident light. Imaginary  $H$  means it radiated in quadrature to change the phase of the light.

To get the microscopic response, we can instead divide the medium of length  $z$  into  $N$  slices. Each slice modifies the field by the factor  $(1 + H/N)$ . Taking the limit of large  $N$  gives

$$\begin{aligned} E_{out} &= \lim_{N \rightarrow \infty} E_0(1 + H/N)^N \\ &= E_0 e^{\kappa z} \end{aligned} \quad (1.35)$$

where  $\kappa = \frac{\partial H}{\partial z}$ . An in-phase or opposite-phase medium re-radiation ( $\kappa$  positive or negative) causes amplification or attenuation, while an in-quadrature response (imaginary  $\kappa$ ) just changes the phase. The slowly varying field envelope propagates according to  $\frac{\partial E}{\partial z} = E_r = \kappa E$ .

It might already be clear that the re-radiation is in quadrature with the oscillation, i.e.  $H \propto iP$ , but we will confirm this in section 1.3.6. This combined with the linear response of the atoms to the field are the two main concepts needed to understand a majority of phenomena concerning our quantum memories. The main situation we haven't covered is when field intensities become large. In our current model the atoms are infinitely polarisable, but in a quantum model they will change state. This is covered by the OBEs.

Although not relevant to our present concerns, it's interesting to note that the physics behind the re-radiation of dipoles is not totally clear, even in the fully quantum case.

For our classical dipole, it must be this way so that the in-quadrature radiated light exerts a damping force on the dipole itself, thus preserving energy. This is the “radiation-reaction” force of Abraham-Lorentz[39]. In this picture the radiation is associated with the dipole’s rate of change of acceleration, rather than the acceleration itself. However this only holds for periodic oscillations and breaks down at the quantum scale. In a quantum picture, the role of the radiation-reaction force can be largely taken up by interactions with the vacuum, although there are cases when the radiation-reaction remains a useful interpretation[40]. Yet even in a quantum model infinities occur, solved for the purposes of calculations only relatively recently (1950s) by renormalization[41]. While a functional solution, renormalization is considered by many to be philosophically unsatisfactory. This is in part because the theory of gravity suffers similar self-force difficulties which cannot be solved in the same way. Thus our problem of dipole re-radiation is related to current efforts in string theory. Maxwell’s equations will be sufficient for our purposes.

### 1.3.4 Optical Bloch Equations

A more general model for an atom is given by the semi-classical OBEs, which we now derive briefly. As two level systems, our atoms may be considered equivalently to spin-half systems or qubits[12], and represented on the Bloch-sphere. They are thus in a sense 3-dimensional, whereas our dipoles are only 2. The 3rd dimension is related to the excitation. Several ways of deriving the OBE exist: one of particular note is by Supplee[42] in which they are found as a simple empirical extension to the Lorentz model to describe quantum effects. The following is similar to the approach by Suter[38] or Allen and Eberly[43].

The density matrix of a single atom<sup>11</sup> (or ensemble of atoms) is

$$\rho = \begin{pmatrix} n & \alpha/2 \\ \alpha^*/2 & 1-n \end{pmatrix} \quad (1.36)$$

where  $\alpha$  is a coherence terms which is proportional to the polarisation in the Lorentz model. The quantity  $n$  describes the average level of excitation. It is not present in the Lorentz model.

The time evolution is governed by  $\dot{\rho} = -[\rho, H]$ . The Hamiltonian,  $H$ , is particularly simple:

$$H = H_{atom} + H_{light} = \begin{pmatrix} 0 & 0 \\ 0 & \hbar\omega_0 \end{pmatrix} + \begin{pmatrix} 0 & g \\ g & 0 \end{pmatrix} \Re(E_0(t)e^{i\omega_0 t}) \quad (1.37)$$

where  $\omega_0$  is the resonant frequency of the transition and  $g$  is its dipole moment.

We get rid of the rapid oscillations by moving into a rotating frame and making the

---

<sup>11</sup>that is an atom that is not necessarily in a pure state

rotating wave approximation. We get for the dynamics of the density matrix:

$$\dot{\rho} = \begin{pmatrix} \Im(g\alpha E) & -i(1-2n)gE/2 \\ i(1-2n)gE^*/2 & -\Im(g\alpha E) \end{pmatrix} \quad (1.38)$$

Separating the matrix elements and letting  $w = 2n - 1$  we get:

$$\dot{\alpha} = -i w g E(t, z) \quad (1.39)$$

$$\dot{w} = g \Im(\alpha E) \quad (1.40)$$

There is no damping and no spontaneous emission here because we are treating the field classically. It will suffice for us to include the decay phenomenologically. We require that  $w$  and  $\alpha$  decay to their equilibrium values exponentially with the time constants  $T_1$  and  $T_2$ .

$$\dot{\alpha} = -\frac{1}{T_2} \alpha(t, z) - i w g E(t, z) \quad (1.41)$$

$$\dot{w} = \Im(\alpha E) - \frac{1}{T_1} (w + 1) \quad (1.42)$$

Describing the relaxation in this way was first considered by Felix Bloch[44].  $T_1$  and  $T_2$  are sometimes known as the longitudinal and transverse decay times. The first corresponds to population decay, so is caused by processes that change the state, such as spontaneous emission. It is a decay rate of energy.  $T_2$  on the other hand gives the decay of coherence, and of polarisation. It can be caused by many mechanisms, including the population decay rate. In the absence of other mechanisms, it is limited to  $T_2 = 2T_1$ . Mechanisms which purely influence the coherence, without affecting the population may be described by another decay time written  $T_2^*$ . The total transverse decay time is then given by  $1/T_2 = 1/2T_1 + 1/T_2^*$ .

We will often describe atoms in terms of a decay rate, rather than a decay time. Further for most situations in this thesis it is sufficient to ignore extra dephasing,  $T_2^*$ , on a single atom (we will definitely need to consider dephasing due to ensemble effects though). In this case we just use  $\gamma = 1/T_1 = 2/T_2$ .

Finally it's common to separate out the real and imaginary parts of the coherence  $\alpha$ . This makes the Bloch-sphere representation particularly clear. Letting  $\alpha = u - iv$  we get

$$\begin{aligned} \dot{u} &= -\Delta v - \frac{\gamma}{2} u + w g \Im(E) \\ \dot{v} &= \Delta u - \frac{\gamma}{2} v + w g \Re(E) \end{aligned} \quad (1.43)$$

where we have also included a detuning term,  $\Delta$ . The  $u$  and  $v$  are analogous to those considered earlier in the Lorentz model.

The factor  $\Omega = g|E|$  is known as the Rabi frequency, and gives the rate at which a resonant atom is excited. More general behaviour of the atoms may be expressed in terms

of the generalised Rabi-frequency,

$$\Omega = (gE, 0, \Delta)|\Omega| = \sqrt{\Omega_0^2 + \Delta^2} \quad (1.44)$$

for short-times compared to the damping.

We note that equations (1.43) reduce to those derived for the Lorentz oscillator (1.33) when  $w \approx -1$ . The two models thus give the same behaviour when the incident light field is small. This is mainly the regime we'll work in.

### 1.3.5 Bloch Sphere Dynamics

The Bloch-sphere is a very useful tool to visualise the evolution of ensembles of two-level systems, such as our atoms. Most dynamical processes correspond to a simple transformation like rotation or decay. We'll use the Bloch-sphere to show the evolution of an inhomogeneously broadened ensemble —one where each atom has a slightly different resonant frequency— under conditions of driving and decay. We neglect any radiation due to the atoms themselves.

The dynamics of an atom in the Bloch-sphere picture may be understood simply as the combination of three effects: An atom's detuning causes it to precess about the  $w$ -axis at that frequency. An external field causes a rotation about an axis in the  $u$ - $v$  plane. The phase of the field determines the angle that the rotation axis makes with the  $u$  axis, while the amplitude of the field determines the rate of rotation (the Rabi frequency,  $\Omega$ ). Spontaneous decay ( $1/T_1$ ) causes an exponential decay toward  $w = -1$  while other dephasing mechanisms ( $T_2^*$ ) causes a decay inward toward the  $w$  axis. The two rotation effects, due to an atoms detuning and an external field, may be combined into the generalised Rabi-frequency (eq. 1.44). That is, the two orthogonal rotations are equivalent to one about a different axis.

The evolution of an ensemble is depicted under six different conditions in Fig. 1.10. For clarity we only show the paths of 3 representative atoms: two oppositely detuned atoms, and one at zero detuning. In addition we define an ensemble vector representing the mean of a larger ensemble of atoms. This gives the net polarisation of the ensemble, and hence the net radiation due to it.

The situations shown are:

- Free evolution. The atoms start in a superposition state, with  $u = 1$ . They precess about the  $w$  axis according to their detuning, while decaying toward  $w = -1$ . The ensemble polarisation approaches zero quicker than the decay as the vectors move out of phase with each other.
- Steady driving. The atoms start in the ground state, before a constant field is applied at zero detuning. Resonant atoms (i.e. zero detuned) are rotated about an axis in the  $u - v$  plane, depending on the phase of the field. The detuned atoms appear

to rotate about a different axis due to the combination of their precession and the driving rotation. The axis of rotation is given by the generalised Rabi vector.

- **Hard pulse.** A very short, strong, resonant pulse is incident. The Rabi frequency is much greater than the spread of detunings in the ensemble so that the atoms don't have time to precess. They thus experience the same rotation. The *area* of the pulse is the Rabi-frequency multiplied by time. This area determines the angle through which the atoms rotate- in this case the area is  $\pi$ . Pulses like this are extremely useful for manipulating atoms, and are named after the angle of rotation. A  $\pi$  pulse may be used create an inversion, or a  $\pi/2$  may be used to create a superposition state(not shown).
- **Inversion by adiabatic passage.** Another method may be used to create an inversion much more slowly, but potentially more robustly. In this example, the 'adiabatic passage' is created by slowly scanning a far off-resonant beam through resonance and past again. The generalised Rabi-vector corresponding to this beam gradually moves from pointing to the ground state to the excited state. The atom Bloch-vectors rotate around the vector and follow it as it moves. The condition for adiabatic passage is that the rate of rotation is much greater than the rate at which the generalised Rabi-vector moves.
- **Small, resonant pulse.** A resonant pulse leaves the atoms excited after it has passed, and thus energy is left in the medium. If each atom were alone in the medium, it would re-radiate most of this energy (unless its decay time is shorter than the length of the pulse). Instead the ensemble polarisation decays rapidly according to its spectral width, and the ensemble stops radiating long before an individual ion would. The energy thus remains in the medium.
- **Small, off-resonant pulse.** An off-resonant pulse causes an adiabatic excitation which is reversed to leave the ensemble in the ground state. The first part of the pulse exciting the atoms and the second part de-exciting them, to leave no energy behind."

The ensemble vector represents the medium as a whole, and its projection on the  $u$ - $v$  plane indicates the average amplitude and phase of radiation of all the individual atoms which make up the ensemble. In many common experiments the vector behaves similarly to that of a single atom<sup>12</sup> with a mean detuning— only with an more rapid decay in the  $u$ - $v$  plane. This rapid decay is qualitatively different from that of the individual ions however: the decay of a single ion according to  $T_2$  represents information lost by interaction with the external environment. The additional decay of the ensemble vector is due to the interaction between ions, and as such no information is lost. Quantum memories are implemented by reviving the ensemble coherence at a later time.

---

<sup>12</sup>more strictly it behaves similarly to a packet of identical atoms, unless we consider the average over many identical measurements

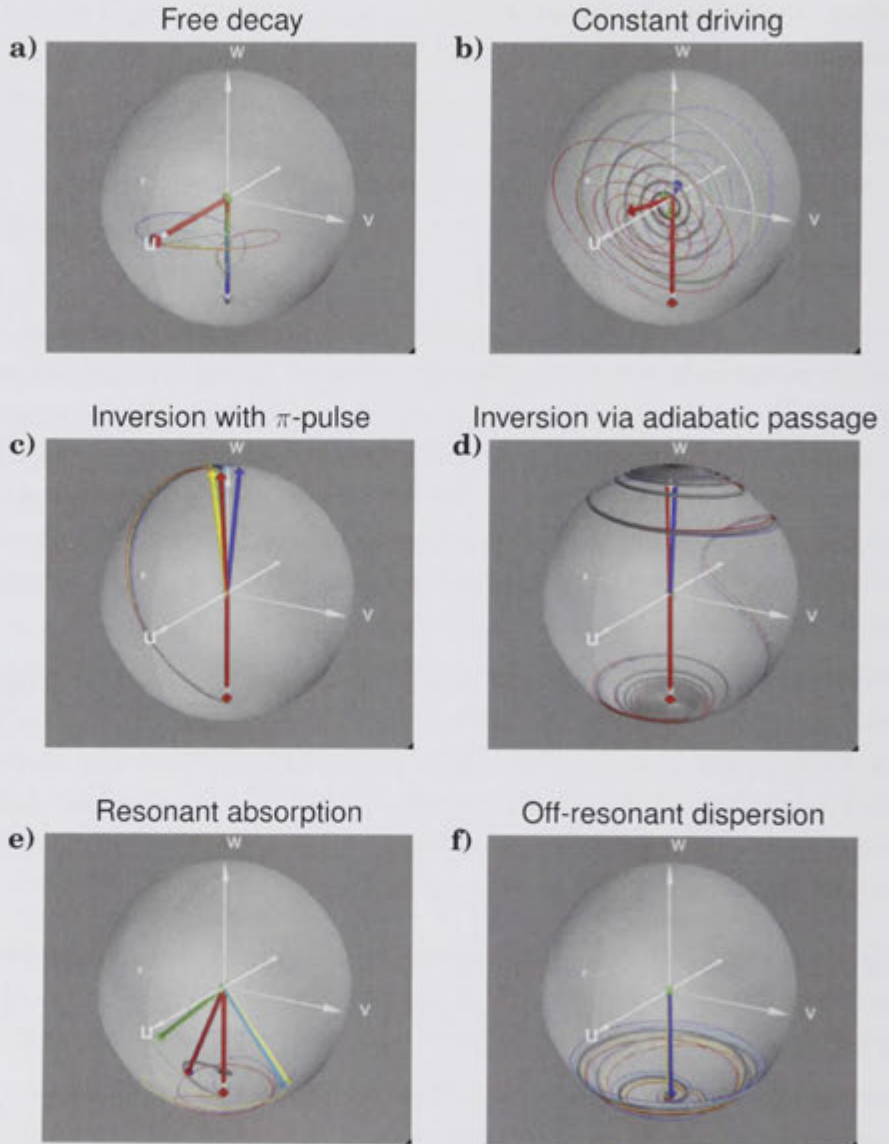


Figure 1.10: The evolution of Bloch-vectors representing an ensemble of atoms. Green vectors are at zero detuning, blue and red are opposite. The grey line represents the ensemble average. For detailed descriptions, see the text. a) Free evolution. b) Steady driving. c) Hard pulses. d) Adiabatic passage. e) Small, resonant pulse. f) Small, off-resonant pulse.

### 1.3.6 Maxwell-Bloch Equations

Combining Maxwell's equations with the OBEs gives us the so called MBEs for describing light propagation in our medium of two-level atoms. These describe a rich array of non-linear phenomena when the light intensity is large, although our main interest is the low intensity, linear regime.

We limit ourselves to Maxwell's wave equation in 1D, written

$$\left[ \frac{\partial^2}{\partial z^2} - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right] E(t, z) = \frac{1}{c^2} \frac{\partial^2}{\partial t^2} P(t, z) \quad (1.45)$$

If  $P = 0$ , this is just the well-known wave equation supporting solutions moving uniformly in each direction with speed  $c$ , such as  $e^{i(\omega t \pm k_0 z)}$ . This combined with an expression for  $P$  from the OBEs gives the so-called MBEs, although in a form that is relatively unwieldy. A common alternate approach is to assume that  $P$  is a function of  $E$ , as:

$$P(t) = \epsilon(\chi^{(1)} E + \chi^{(2)} E^2 + \chi^{(3)} E^3 \dots) \quad (1.46)$$

In the low-excitation regime, as in the dipole picture, the response is linear and only the first term is needed. Inserting  $P = \epsilon_0 \chi E$  gives

$$\left[ \frac{\partial^2}{\partial z^2} - \frac{1 + \chi(\omega)}{c^2} \frac{\partial^2}{\partial t^2} \right] E(t, z) = 0 \quad (1.47)$$

So that the effect of the linear interaction is just a change of refractive index, giving  $n = \sqrt{1 + \chi(\omega)}$ . The new wave solutions that are supported contain the physics of propagation. They are  $E = E_0 e^{i(\omega t \pm k z)}$ , with the new wave vector  $k$  being frequency dependent according to the dispersion relation,

$$k^2 = (1 + \chi(\omega)) k_0^2 \quad (1.48)$$

$k_0$  is the vacuum wave-vector. This is how the medium's polarisation in response to the field, as described by  $\chi$ , determines light propagation in the medium. An imaginary  $\chi$  causes attenuation, a real  $\chi$  causes a phase shift (dispersion).

By comparing this wave-vector with the propagation found 1.3.3, we obtain for the 'radiation' response

$$\kappa = ik_0(1 - \sqrt{1 + \chi}) \quad (1.49)$$

, i.e.  $\kappa$  is the difference between the vacuum and dielectric wave-vectors. It's the part which represents the effect of the medium. If  $\chi$  is small, we can approximate  $\kappa \approx -ik_0 \chi / 2$ , while if  $\chi$  is large,  $\kappa \approx -ik_0 \sqrt{\chi}$ . If our SVEA is to hold, the new wave-vector must be similar to the old, and  $\kappa \ll k_0$ . We will assume from here on that  $\chi \ll 1$  and that  $\kappa = -ik_0 \chi / 2$ , and note that for the situations we are concerned with,  $\chi, \kappa \approx 0.001$ . In

this approximation it is clear how our dipoles re-radiate. Given  $P = \epsilon_0 \chi E$ ,  $E_r = \kappa E$ ,

$$E_r = \frac{\kappa}{\epsilon_0 \chi} P = -i \frac{k_0}{2\epsilon_0} P \quad (1.50)$$

That is, the radiation lags  $\pi/2$  behind the polarisation, as we might have guessed. This forms the basis of a very useful, simplified form of the MBEs. We will see this more clearly below.

### Approximate form of the Maxwell-Bloch Equations

As long as we're in the linear (small excitation) regime, and the medium doesn't change with time (or space), we only need the new wave-vector,  $k$  (or  $\chi$  or  $\kappa$ ) to describe all effects. In our SVEA approximation however, the Maxwell and Bloch equations can be made quite tractable to solve more general problems. The resultant expressions are also suitable for numerical simulations as all rapid variations are removed. We might arrive at them by pushing our conceptual approach in 1.3.3 further, but the more common approach starts with Maxwell's equation in 1D (1.45). Detailed steps are given in many places, e.g. Allen and Eberly [43].

A general solution of eq. 1.45 in the SVEA is  $E_{right}(t, z)e^{i\omega t - kz} + E_{left}(t, z)e^{i\omega t + kz}$ . However we only need consider pulses in one direction: as part of the SVEA we assume there are no reflections as would be caused by abrupt changes in space, and thus a forward propagating wave stays forward propagating. We also assume a slowly varying envelope for  $P$ . Substituting in 1.45 and neglecting small quantities we obtain:

$$\left(\frac{\partial}{\partial z} - \frac{1}{c} \frac{\partial}{\partial t}\right) E(z, t) = -i \frac{k_0}{2\epsilon_0} P(z, t) \quad (1.51)$$

The two derivatives can be combined if we make a transformation to coordinates moving at the speed of light  $t_c = t$ ,  $z_c = z + ct$ , such that  $\frac{\partial}{\partial z_c} = \frac{\partial}{\partial z} - \frac{1}{c} \frac{\partial}{\partial t}$ . This is often called the co-moving frame. Finally we have the remarkable simple result

$$\frac{\partial E}{\partial z_c} = -i \frac{k_0}{2\epsilon_0} P \quad (1.52)$$

which of course is just  $E_r$  as above, the slowly-varying envelope of the radiation due to polarisation density  $P$ . This combined with the OBEs (1.43) for the evolution of the polarisation comprises the MBEs for slowly-varying envelopes.

### 1.3.7 Spectral Response, Absorption and Dispersion

We've already discussed a few ways to describe the interaction between a linear medium and a light field—the susceptibility,  $\chi$ , the radiant response,  $\kappa$ , and the dispersion relation governing the wave-vector  $k$ . The most directly useful when discussing the effect of the medium on the light is  $\kappa$ , and we emphasise it here. Other useful factors that may be



written in terms of  $\kappa$  are the total transmission, logarithmic absorption  $A$ , and the group velocity/delay. These are plotted for a single damped oscillator in Fig. 1.11.

The spectrum of  $\kappa(\omega)$  for the damped oscillator is given in Fig. 1.11a. If light with spectrum  $E(\omega)$  is incident on a narrow slice of medium with thickness  $\delta z$ , the medium will radiate in response as  $E_r = E(\omega)\kappa(\omega)\delta z$ . Thus Fig. 1.11 directly describes the relative change of the field envelope with distance. Negative, real  $\kappa$  describes atoms re-radiating out of phase causing attenuation. Imaginary  $\kappa$  describes re-radiation in-quadrature causing a change of phase.

The effect of the medium as a whole may be described by the absorption and transmission of the field through the medium. For a uniform medium of length  $L$ , the field transmitted is given by  $T = e^{\kappa L}$ . This plotted in Fig. 1.11c. The fraction of energy absorbed by the medium is  $A = e^{2\kappa_r L}$ . The absorption is particularly convenient as a logarithm, in which case we only need the exponent. In units of absorption lengths we have  $A_n = -2\kappa_r z$  (this is often called “ $\alpha L$ ” in reference to Beer’s law with attenuation coefficient  $\alpha$ ) or decibels  $A_{dB} = -20 \log_{10}(e)\kappa_r z$ . This means that combining the effect of additional absorbers or extra medium length is a simple addition. Even more usefully, the shape of a logarithmic absorption spectrum is independent of its height. This allows consistent definitions of spectral line-widths and shapes, regardless of the length of the medium. The logarithmic absorption is plotted as a function of frequency in Fig. 1.11b.

### 1.3.8 Slow Light

Interaction with the medium may change both the phase velocity and the group velocity of light travelling in it. A change in overall phase velocity (and an equal change in group velocity) is responsible for the refractive index of most transparent materials. The phase velocity is the speed of a peak (or trough) of a monochromatic wave.

Near resonance, the group velocity can be drastically different from the phase velocity. The group velocity is the speed at which a pulse, or wave-packet progresses, and is given by the rate at which the phase velocity changes with frequency. We may think of a pulse as a sum of many spectral components. The peak of the pulse occurs when all the components are in phase. We can express the condition for this to hold as  $\frac{\partial \theta}{\partial \omega} = 0$  where the phase is given by  $\theta = \omega t - k_0 z + \kappa_i z$ , where  $\kappa_i$  is the imaginary part of  $\kappa$ . We get

$$t - z \left( \frac{\partial k_0}{\partial \omega} + \frac{\partial \kappa_i}{\partial \omega} \right) = 0 \quad (1.53)$$

We then have

$$v_g = \frac{z}{t} = \frac{c}{1 + c \frac{\partial \kappa_i}{\partial \omega}} \quad (1.54)$$

where  $c = \frac{\partial k_0}{\partial \omega}$  is of course the speed of light in vacuum. It is mathematically simpler in our approximations to work with the reciprocal group-velocity, the group-delay per unit length. This is

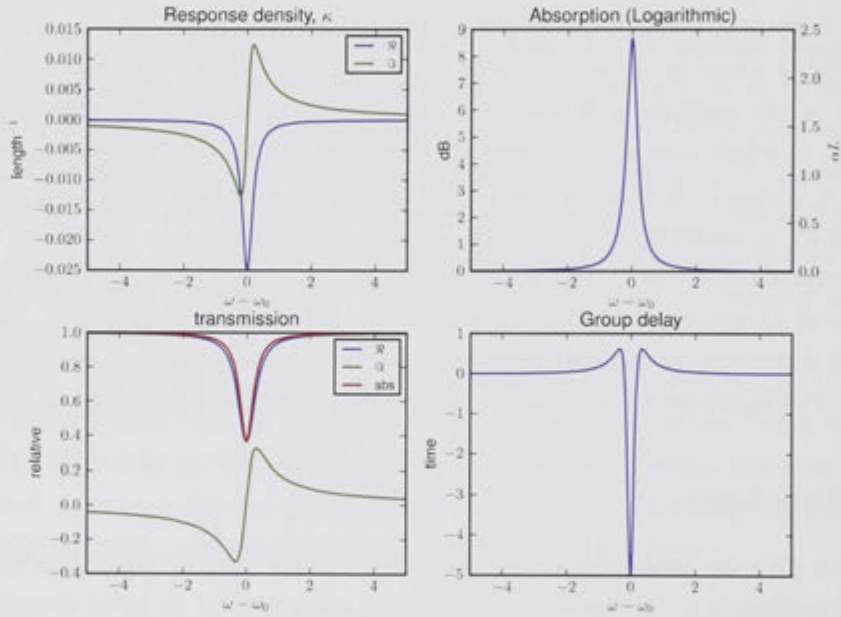


Figure 1.11: Spectral properties of a medium with a single, weakly damped resonance. a) The response density  $\kappa$ . b) Logarithmic absorption (of energy). Left axis is dB, right axis is absorption lengths ( $\alpha L$ ). The FWHM maximum of the line is given by  $\Gamma = \frac{1}{\pi T_2} = \frac{\gamma}{2\pi}$ . c) Transmission, in terms of field. d) Group delay of the medium. Near, but not on resonance, there is a positive group delay meaning a pulse within this bandwidth is effectively slowed. On resonance the group delay is negative, apparently speeding up a pulse in this window. This “fast-light” however only occurs in the presence of strong absorption and doesn’t speed up information—in some sense the peak is shifted forward in time simply by removing the part that arrives later.

$$\frac{1}{v_g} = 1/c + \frac{\partial \kappa_i}{\partial \omega} \quad (1.55)$$

The second term is the delay in our SVEA, and is the total delay in a reference frame moving at the speed of light.

To calculate the total delay caused by the medium we simply integrate over  $1/v_g$ , as  $c \int \frac{\partial \kappa_i}{\partial \omega} dz$  (or  $= c \frac{\partial \kappa_i}{\partial \omega} z$  for uniform media).

The group velocity and group delay are shown in Fig. 1.11 for a homogeneously broadened transition (or a damped dipole oscillator).

### 1.3.9 Population Inversion and Amplification

Our assumptions of linear response for polarisation and re-radiated field require that the excitation remain low. The linear assumption holds so long as energy is proportional to the square of the polarisation— i.e. that the bottom of the Bloch-sphere approximates a parabolic shape. This is clearly not the case when absorbers are excited to near the equator. When the atoms are fully excited however, the response again becomes linear. This time the response is of opposite sign. The atoms will respond to incident light by radiating *in-phase*, causing amplification instead of attenuation. This response reversal is the basis of the ‘photon-echo’. It is also the reason it cannot be used as a quantum memory, as the amplification introduces noise.

### 1.3.10 Impulse Response

The inverse Fourier transform of the medium’s spectral response,  $H(\omega)$ , is  $h(t)$ . This is known as the impulse response of the medium<sup>13</sup>. While  $H$  gives the response to a delta-function in frequency space,  $h$  gives the response for a delta function in time. One property of  $h(t)$  is that it must obey causality: the requirement that no response occurs before a signal is input. The equivalent restriction on  $H(\omega)$  means it is completely defined by either its real or imaginary components (absorptive or dispersive character). The relation between the components is known as the Kramers-Kroenig relations[45].

The impulse response can be convenient to look at phenomena in the time domain. One simply treats each point of an input pulse as an impulse. Thus output is the convolution of the input pulse with  $h(t)$ .

The impulse response of an inhomogeneously broadened ensemble of atoms is of particular note. It’s decay corresponds to the dephasing as seen on the Bloch-sphere in Fig. 1.10. While it appears the same as the decay of a single absorber with a short  $T_1$ , for the inhomogeneous decay the energy remains in the system.

If the medium is altered in some way during the interaction time, none of these descriptions remain strictly valid. Yet time dependence is clearly a requirement of memories.

<sup>13</sup>It is closely related to the FID, which usually describes the transient response after a long pulse is switched off suddenly

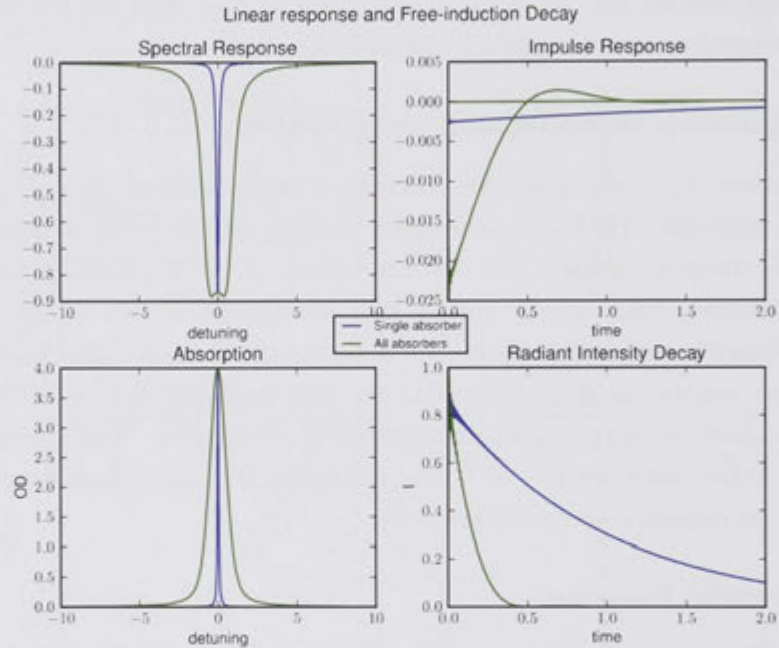


Figure 1.12: Entire medium response functions and free-induction decay from a short pulse. We consider a medium consisting of weakly damped atoms (the normalised behaviour for a single atom is represented by the blue lines) which have a Gaussian distribution of detuning frequencies. This causes the medium as a whole to appear as if it is a single strongly damped atom (green lines). The impulse response (no units, top right) describes the response of the medium when hit with a short sharp pulse: it absorbs light, then re-radiates it after it has passed. The re-radiated intensity at a given time is related to the absolute value of the impulse response squared (bottom right). The decay rate for the single atoms is often due to damping losses. The shorter decay for the broadened medium however is due to the radiation of the constituent atoms interfering destructively, while the energy remains stored in the microscopic polarisation. The energy loss rate due to damping remains the same as for an individual atom.

The above descriptions remain useful so long as a medium behaves consistently during the time of the impulse response.

### 1.3.11 Time-dependant Response: Echoes

Thinking in terms of the medium's response, absorption occurs because atoms configure themselves to radiate out of phase with an incident pulse. The response adds destructively with the pulse so the net forward-going wave is zero. As we have seen, the absorption is not necessarily dissipative if it's on a shorter time scale than  $T_2$  of a single atom. A photon echoes can occur when the medium is coaxed to repeat the absorption configuration after the input pulse is passed. This causes a re-emission, or 'echo' as this time the medium radiation has no input pulse to destructively interfere with it. There are several ways to cause this re-coherence which we'll discuss in the next chapter. The purpose of this section is to give a rough and hopefully intuitive description of a general echo.

We have already discussed something which looks like a memory: slow-light. As described earlier, off-resonant (dispersive) interaction can cause the delay, or temporary absorption, of a pulse. However in the systems we've discussed so far, the delay time is fixed and cannot be used to recall the pulse on demand. To implement an actual 'memory', the medium response must depend on time. In principle this greatly complicates the problem: to fully define a time-dependent medium, an impulse response  $h(t, t_0)$  must be given for every point in time  $t_0$ . We wish to consider only a simple class of the general problem.

For most situations of interest in this thesis, the response of the medium only changes once, and this is after the pulse has been absorbed. Thus the medium may be considered time-independent while the pulse is incident. A further simplification is to assume that the medium response is fast compared to the length of the pulse (equivalently that the medium absorption bandwidth is greater than that of the pulse). After the pulse has been absorbed and the medium is no longer responding, it is modified to cause the re-emission. The details of this we'll leave until chapter 2. A simple echo is described in Fig. 1.13.

In the input stage, the medium responds to the pulse by configuring to radiate a field  $E_r = \int_{-\infty}^{\infty} E_{in} h(t) dt$ . If the response of the medium is quick enough,  $E_r$  is just a scaled, inverted copy of the input so that  $E_r = s E_{in}$  ( $s$  negative) (a short enough  $h(t)$  approximates a negative, scaled  $\delta$  function). This is the red trace in Fig. 1.13.

If the medium response is strong (i.e. optical-depth of the medium is large), the light is completely absorbed. The medium response is  $E_r = -E_{in}$  ( $s = -1$ ). Considering the atomic evolution, the light drives them to an excited state where they quickly move out of phase with each other according to their detunings. The individual atoms would radiate but the ensemble average is zero. An echo is achieved by reversing the relative evolution of the absorbers: in the diagram this is done by reversing their detunings. As they rephase, they recreate the input field exactly, but with opposite phase. Thus while the input field drove them up on the Bloch-sphere, the field they create drives them down again. The

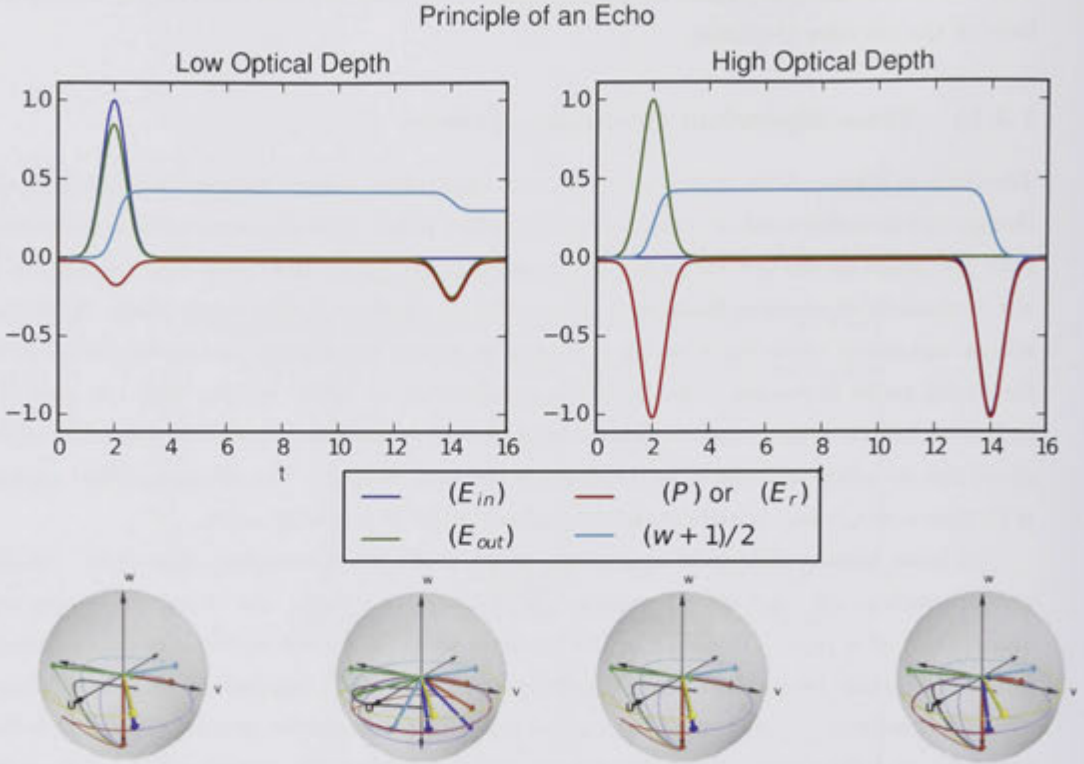


Figure 1.13: General principle of a photon echo. To absorb light, a medium reacts to a pulse by radiating out of phase, cancelling it. The principle of an echo is to re-create the situation of absorption- only this time without the input light to cancel. When coupling is high, the process is exactly reversible, and the echo drives the atoms to the ground state just as the input pulse drove them up. When atom-field coupling is low, the absorption is not exactly reversed and some of the excitation remains on the atoms.

situation seen by the atoms then is a perfect time reversal. They finish in the ground state and a time-inverted input pulse is re-created (only in this case with opposite phase).

The weak response (optically thin) case is identical except for the scale of the response. During the input, only a fraction of the light energy is transferred to the atoms. Then during the re-emission, the field the atoms see is smaller than that during the absorption. The atoms thus do not drive themselves back to the ground state, and thus only re-emit a fraction of the energy they absorbed. In terms of field however, the output is actually larger than  $E_r$  by a factor of 2. This is because without the de-excitation as in the high OD case, the mean coherence is twice as high as it would be otherwise. This may be proven by more rigorous analysis, but is merely stated here.

The above describes roughly most ‘echo’ techniques currently considered, with one exception: we are ignoring propagation effects due to the finite size of a medium. This is as if we are using an infinitely thin, but very optically thick slice. This is the complete

opposite of the SVEA, and thus our previous analysis doesn't strictly hold. If we allow propagation effects, in the optically thick case the echo will simply be absorbed by a later part of the medium. In practise, a good echo is made if the output step 'looks' identical to the input, spatially as well as temporally.

The direct way to achieve this is to reverse the direction of propagation at the same time as the time evolution switches. Reversing the propagation direction requires reversing the spatial phase (described by  $e^{ikz}$  for light of wave-vector  $k$ ) of the created coherence. This can be achieved using an additional ground state, or a pair of  $\pi$  pulses at specific angles. We now describe this briefly.

### 1.3.12 Transferring Coherence and Changing the Emission Direction

Mostly we are concerned with 2-level systems, however additional levels provide more degrees of freedom, and thus many more possibilities in the kinds of interactions that can occur. In particular an extra level can be used to store coherence for a long time, or to change the direction of an echo. We consider here addressing only a single transition at one time. More general interactions where more than two levels are addressed simultaneously are also relevant but are better covered with a dressed atom picture (see e.g. Compagno et al. [46]).

If we ignore propagation effects the process is straight-forward. We consider a pure state of a  $\lambda$  system. The two ground states are very close in frequency so that the wavelength between each and the excited state is the essentially the same. The process is well described by tracking the coherence of each transition in a  $3 \times 3$  density matrix (see e.g. Suter [38]), but the physics may be clearer if we track the phase of each level. The initial state is

$$|\psi\rangle = 1|g1\rangle + 0|e\rangle + 0|g2\rangle \quad (1.56)$$

After some excitation by a pulse travelling in the  $z$  direction, some coherence is created whose phase depends on the phase of the light which created it. The result is

$$|\psi\rangle = \sqrt{(1-\epsilon)}|g1\rangle + \sqrt{\epsilon}e^{ikz}|e\rangle + 0|g2\rangle \quad (1.57)$$

where epsilon is determined by the amplitude of the pulse.

We apply a  $\pi$ -pulse again in the  $z$  direction, but this time resonant with the  $|e\rangle - |g2\rangle$  transition. This is essentially the reverse process of the input pulse. It will swap the populations of the levels, and in the process, it will undo the spatial phase that was created (the rotation each vector sees is  $e^{i(\pi-kz)} = e^{ikz}$ ). However there is now a definite phase relationship (a coherence) between the ground state levels

$$|\psi\rangle = \sqrt{1-\epsilon}|g1\rangle + 0|e\rangle + \sqrt{\epsilon}|g2\rangle \quad (1.58)$$

As these are much lower energy transitions, e.g. spin levels, the coherence may poten-

tially last much longer.

Finally, the coherence may be transferred back to the excited state for readout by another  $\pi$  pulse on the  $|e\rangle - |g2\rangle$  transition. The direction of this pulse determines the spatial phase of the coherence created. Thus if the  $\pi$  pulse is applied backward, radiation due to this coherence will propagate backward

$$|\psi\rangle = \sqrt{1-\epsilon}|g1\rangle + \sqrt{\epsilon}e^{-ikz}|e\rangle + 0|g2\rangle \quad (1.59)$$

This direction reversal is significant in an optically thick sample as it may allow the re-emission of light backward, that is without re-absorption by later parts of the medium.

The slightly more general case may be expressed as phase-matching the  $k$ -vectors of the pulses such that:

$$\vec{k}_{coh} = \vec{k}_{in} - \vec{k}_{\pi1} + \vec{k}_{\pi2} \quad (1.60)$$

where  $k_{coh}$  is the wave-vector of the created coherence.

It is possible to change the spatial coherence without using an additional level, using 2  $\pi$  pulses on the same transition. The phase-matching condition is different in this case, as each  $\pi$  pulse alters the phase of population in both levels at a time. This doubles the effect of each pulse, giving us for the phase-matching condition:

$$\vec{k}_{coh} = \vec{k}_{in} - 2\vec{k}_{\pi1} + 2\vec{k}_{\pi2} \quad (1.61)$$

The condition for  $\vec{k}_{coh} = -\vec{k}_{in}$  occurs if  $\vec{k}_{in} - \vec{k}_{\pi1} + \vec{k}_{\pi2} = 0$ , which is matched if the pulses are in a cross pattern from the side, with an angle of  $60^\circ$  between the input direction and that of the first  $\pi$  pulse, then a further  $60^\circ$  to the second. Thus it is still possible to reverse the propagation direction with only 2-levels available, but it's experimentally trickier.

Phase-matching of two level echoes like this was first discussed in 1966 by Abella et al. [47].





---

# Coherent and Quantum Memories in Atomic Ensembles

---

The term “coherent” in quantum mechanics implies that phase is preserved. For our purposes we’ll apply the term a little more strictly: a coherent memory is one that preserves information about the input’s complex amplitude. This means that information both about its phase and magnitude are preserved<sup>1</sup>. A quantum memory may then be defined as a very precise coherent memory: precise enough that quantum information is preserved which can be used for quantum information applications.

To make the definition concrete, we need to precisely define quantum information. Doing this generally can be a little tricky<sup>2</sup>, but a practical definition is easy: quantum information is that which is lost when we make a measurement. The border between a memory that is ‘quantum’ and one that is merely coherent is given by the best performance of a device based on measurement- a “classical” memory. A quantum memory must be coherent to be considered quantum, because it is the simultaneous measurement of conjugate variables (i.e. those related by phase) that is limited by quantum mechanics.

The major difficulty in implementing quantum memory for light is obtaining strong coupling between light and a quantum system, while keeping that system very weakly coupled to the rest of the universe. A powerful way to achieve this is to use weakly interacting atoms, but lots of them: an ensemble.

Such ensembles have been studied for many decades for the coherent, but ‘classical’, storage and manipulation of light. It is only in recent years that their use for ‘quantum’ storage has been targeted. The earlier research is very transferable, as the desirable properties of classical and quantum storage are very similar. In the quantum case, the requirements are simply more stringent.

The present chapter concerns properties and issues of high performance coherent memories using ensembles, and their use in quantum applications. It may be logically divided into 4 parts.

---

<sup>1</sup>Note that for this to be true the output amplitude is not necessarily the same as the input, we only require that there is some one-to-one relationship such that knowing what is output, the input signal can be deduced

<sup>2</sup>see e.g. Brukner[48] and Griffiths[49] for some discussion.

The need for a quantum memory is related to the quantum limitations of measurement and copying. Our first task is to clarify what this means by considering the limitations of simple devices based on measurement for storing, copying, and amplifying quantum states.

Next, we identify general parameters of a coherent memory, and briefly survey specific techniques for realising one. These techniques are all based on the same principle of the ensemble's linear response to light as discussed in section 1.3. They thus have the same fundamental limitations, determined by the spectral density and coherence of the atoms used. We will look at when these limits may be realised.

Having described the performance of coherent memories, we look more closely at the parameter particularly important for quantum applications: noise. We answer the question of why a device which loses 99.9% of the signal input into it can still be considered for high precision quantum applications. Some benchmarks for experimentally determining performance are discussed.

Finally we look at the application area of quantum cryptography, making some statement about what level of performance is required of a memory. We comment on requirements of the storage medium for achieving this performance.

## 2.1 Classical Memories, Quantum Cloners, Amplifiers, and Measurement

The border between 'classical' and 'quantum' domains for our purposes is that which is possible using measurement. To define this border then, we will consider simple implementations of devices operating at the limits of measurement. Apart from a 'classical' memory, we also consider devices to copy, and to amplify quantum states. These will serve to illustrate the problem solved by a quantum memory, and will be used in 2.7 to define more concrete limits. We describe the general devices first, before the specific implementations in Fig. 2.1. To simplify the discussion, we will limit ourselves to symmetrical, or "phase-insensitive" devices that treat all phases equally<sup>3</sup>.

Simultaneous measurements of position and momentum were discussed by Arthurs and Kelly[50]. These correspond to measuring orthogonal phase quadratures of our light field. In principle, a single quadrature could be measured exactly, at the expense of completely disrupting the other. A 'classical' memory is one which attempts to measure both, then recreates the state as accurately as possible when it is required. A phase-insensitive aims for an equally precise measurement of both quadratures.

By making a measurement, a classical memory creates a permanent record of information about the input. In principle, this may be used to create arbitrarily many identical outputs. In this sense, a classical memory may also be called a  $1 \rightarrow \infty$  copying, or 'cloning'

---

<sup>3</sup>The label "phase-insensitive" can be misleading. These devices actually aim to be as sensitive as possible to phase, just not any particular phase

device. In comparison, an ideal quantum memory makes no measurement of the input and thus no record of the state remains. We might call this a  $1 \rightarrow 1$  cloner.

A  $1 \rightarrow N$  ‘cloning’ device gives a natural generalisation to a memory. A value for  $N$  between 1 and  $\infty$  corresponds to making a partial measurement of the input state, leaving behind only partial information about the outputs. The recorded information is by definition ‘classical’, any quantum uncertainty having been lost. The quantum uncertainty that does remain is divided evenly among the  $N$  outputs.

The performance of a cloner when  $N = 2$  gives an important benchmark for other devices, ‘no-cloning’ limit[51]. If the output of a device (e.g. a memory) is a better copy of the input than the best output of a  $1 \rightarrow 2$  cloner, it represents more quantum information than was left behind. A general definition of the no-cloning limit depends on a precise definition of ‘better’, and is the subject of 2.7. Limits of quantum copying were first discussed by Wootters and Zurek[52].

The final process considered is amplification. This is closely related to storing and copying. The quantum limits of amplification were first discussed by Caves[53]. We’ll use an implementation of an amplifier as a bridging step between a simple quantum memory and a more general cloner. The limitations of amplification will also aid in our later discussion of loss (section 2.7).

The limitations of the 3 processes of copying, amplifying, and classically storing quantum states all come back to the limits of measurement. This can be demonstrated explicitly by implementing each device using only linear optics and measurement. Such implementations are given in Fig.2.1. Each device must introduce some level of uncertainty to its output, due to the measurement. For phase-insensitive devices, this may be expressed simply as an added noise (the same over all phases, or quadratures) with variance  $V_{add}$ . We now briefly describe each device.

The classical memory (Fig. 2.1a) simply measures conjugate amplitude quadratures of an input using balanced heterodyne (section 1.2.8). In principle this can meet the Heisenberg limit. The measurement is phase-insensitive, or ‘symmetric’ as we are attempting to measure both quadratures equally<sup>4</sup>. As in section 1.2.8, 2 units of vacuum noise are introduced on each quadrature.

$$V^{add} = 2$$

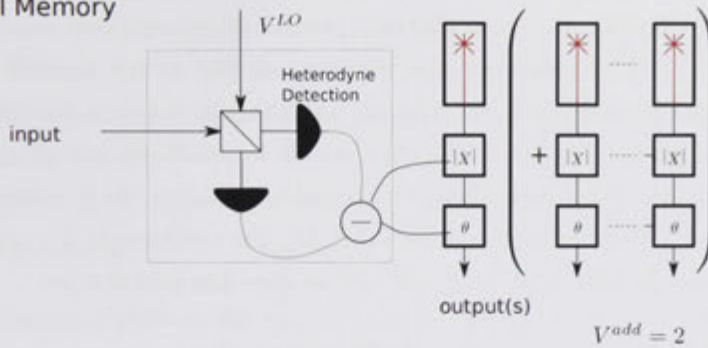
This is sometimes described as paying two quantum duties, or ‘quduties’ for the measurement of the two variables[54].

An optimal quantum-limited amplifier (for Gaussian states) may be realised as a slight modification to our classical memory. Instead of measuring the full input, the device now only measures part of it as picked off by a beam-splitter. The measurement result is used to modify the unmeasured part of the state (Fig. 2.1b), as demonstrated by Josse et al.[55].


---

<sup>4</sup>A general asymmetric device could theoretically be realised by a double-homodyne measurement and a squeezed local oscillator

## a) Classical Memory



## b) Symmetric Amplifier



## c) Symmetric Cloner

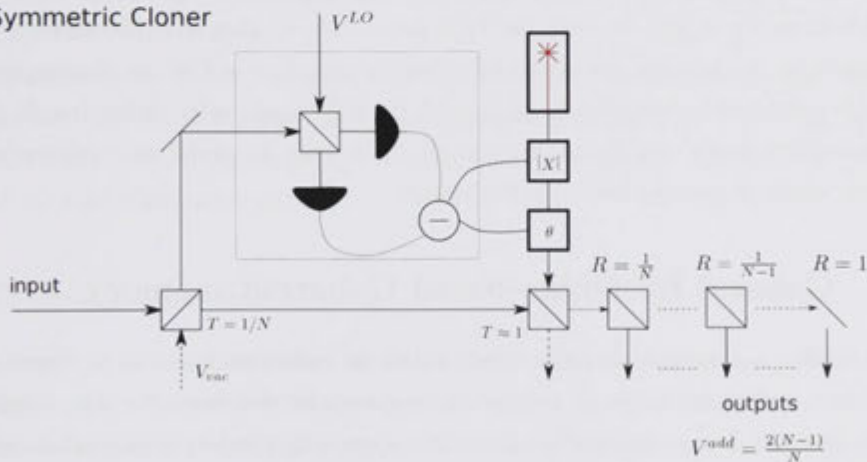


Figure 2.1: Simple implementations of optimal, phase-insensitive, a) ‘classical’ memory, b) amplification, and c) copying, or ‘cloning’. The output of each device must have some classical uncertainty added to it- for our phase-insensitive devices, both quadratures suffer added variance of  $V^{add}$ . The minimum of  $V^{add}$  in each case can be thought of as what strength of measurement is required, as this determines how much of the local oscillator (and its associated noise,  $V^{LO}$ ) is introduced. Note that the vacuum noise introduced at the first beam-splitter in b) and c) does not appear at the output, as it is cancelled by the feed-back[55].

This illustrates that to effectively amplify a state, we need to first know its amplitude. General amplification is then a two step process of measurement, and feed-back. A full measurement will provide the most information, but at the expense of destroying the original state. If a large gain is required, a full measurement is the optimum solution. If we require a gain of 1 however, the state should be unaltered and no measurement made.

The strength of the measurement is altered by changing the transmissivity of a pick-off beam splitter. As described by Josse et al.[55], the optimum is  $T = 1/G$  where  $G$  is the required gain. As in section 1.2.6, the output state has added noise

$$V^{add} > 2(G - 1)$$

A simple cloning device is realised by dividing the output of an amplifier. The gain of the amplifier is set to the number of copies required, that is  $G = N$ . The amplified output is split into  $N$  identical copies using a network of beam-splitters. The noise added by the amplifier is divided by  $N$  to give outputs with

$$V^{add} = \frac{2(N - 1)}{N}$$

For  $N = 2$  we have a  $1 \rightarrow 2$  cloning device, and  $V^{add} = 1$ . This single ‘quduty’ is necessary to preserve the Heisenberg limit- without it we could measure both outputs to learn more about the original[51]<sup>5</sup>.

We note that in this discussion we are aiming to produce outputs with the same mean amplitude as the inputs, so that the total gain of the system is 1. If instead we require a net gain of  $g$ , we may set the amplifier to have a gain  $G = gN$  in our cloning device. This may be considered a slightly more general limit for ‘optimally’ achieving  $N$  attenuated outputs (for certain definitions of optimal). This may be useful as a comparison to real devices which in general have non-unity gain.

## 2.2 General Ensemble-based Coherent memory

We consider a coherent memory which takes an unknown pulse as an input, and may recall it ‘on demand’- that is the recall time may be decided *after* the pulse has been input. Related devices we don’t consider are a one-way memory which takes an input but cannot re-release it as light, a delay line which can only delay a pulse by a fixed amount, and a non-linear memory which cannot take an arbitrary input.

Several techniques for implementing a coherent memory are summarised in Fig. 2.2. We analyse these techniques in the next section, but will first describe the common properties here. These properties are general to a linear, coherent memory.

---

<sup>5</sup>Interestingly, we may further generalise the device to an  $M \rightarrow N$  cloner, that is taking  $M$  identical inputs instead of just 1 [56]. This is done the same as here, except the  $M$  inputs are first concentrated into a single mode (leaving the others as vacuum). This mode is then amplified and split into  $N$  as here.

The usefulness of a memory, for both classical and quantum applications, is broadly determined by its bandwidth, storage time, and mode-capacity. These properties are defined in terms of the efficiency with which information is retrieved.

A fifth property, noise, is the major determinant of whether a memory can be used for quantum applications.

### 2.2.1 Total Efficiency

The efficiency is defined simply as the ratio of energy out to energy in. So long as none of retrieved energy is associated with additional noise (as is the case for amplification), the efficiency is a measure of information preservation. It then corresponds to a quantum efficiency—in a photon picture, it describes the probability of successfully storing and retrieving a photon.

The efficiency in a particular experiment will depend on the length of storage, the properties of the pulse, and the properties of the medium. An overall efficiency may be loosely broken down into factors as

$$\eta = \eta_c^2 \eta_B \eta_T \eta_{bg} \quad (2.1)$$

In order,  $\eta_c$  is determined by the strength of atom-light coupling,  $\eta_B$  is the matching of the pulse bandwidth with the memory,  $\eta_T$  is a decay due to long storage, and the final factor  $\eta_{bg}$  is due to background attenuation and unwanted re-absorption. We define these factors in the following sections.

The breakdown is one of convenience, and the usual time-frequency uncertainty caveats apply. Further, we could consider that the bandwidth and time factors ( $\eta_B$  and  $\eta_T$ ) represent the variation of the coupling efficiency ( $\eta_c$ ) in frequency and time. It remains convenient to treat them separately.

### 2.2.2 Coupling Efficiency

The peak efficiency is first determined by the strength of interaction between the memory and the light. For an ensemble based approach this is determined by the number of atoms, and the interaction strength of each. Often this is the same as the absorption, i.e.  $\eta_c = 1 - e^{-\beta}$  where  $\beta$  (often written as  $= \alpha L$ ) is the optical depth

The factor comes in twice. During the input, it governs how much light is actually absorbed by the medium. Then during the output it governs the efficiency with which the excitation is converted back into light. The total factor for coupling efficiency is then  $\eta_c^2$ .

We must be careful about time scales; we usually give the optical depth in the frequency domain, which by definition assumes infinite time scales. However on time-scales short compared to the inverse of their line-width, atoms contribute off-resonantly. At short time scales of  $\sim \Delta t$ , they interact over a bandwidth  $\sim 1/\Delta t$ .

### 2.2.3 Bandwidth

The memory bandwidth is the spectral width over which the memory operates efficiently. It indicates the reciprocal of the shortest pulse that can be stored, or alternatively the number of longer pulses separated in frequencies. A detailed description of the bandwidth is given by the coupling efficiency as a function of detuning- with the same caveats for time scale. A less precise but more convenient description is the FWHM of this function which we write  $\mathcal{B}$ . If the spectrum of an input pulse is outside of  $\mathcal{B}$ , some spectral components will not be stored efficiently, giving a bandwidth efficiency factor  $\eta_B$  significantly less than 1.

### 2.2.4 Decay Time and Shelving Time

As a quantum state interacts with an environment, the information it contains leaks out (or decoherence leaks in). Thus the efficiency reduces the longer the state is stored. The storage time  $\mathcal{T}$  is fundamentally limited by the dephasing rate of a single member of the ensemble,  $\gamma$ , but is typically quicker due to dephasing between members. In either case, the temporal decay profile is often exponential such that  $\eta_T = e^{-\gamma t}$ . In this case it is convenient to express the storage time as that when the efficiency has decayed to  $1/e$  of the maximum, so  $\mathcal{T} = 1/\gamma$ .

There may be two relevant decay scales for a memory. We usually mean an active period, during which the memory continues to operate and accept more inputs. This time gives the longest pulse, or length of the longest pulse train, that a memory can store

The other possibility is a shelving time,  $\mathcal{T}_{sh}$ . Here the information is already stored in the memory, but it is 'off' in that it can't accept any further input. This is usually implemented by transferring the coherence to an auxiliary ground state. This can be used to dramatically increase the storage time, but not the amount of information that can be stored.

### 2.2.5 Time-bandwidth Product and Mode Capacity

The product of storage time and bandwidth ( $\mathcal{T}\mathcal{B}$ ) gives a dimensionless number that is often used as an indicator of memory performance. Depending on the purpose considered, either the active storage time ( $\mathcal{T}$ ) or the shelving time ( $\mathcal{T}_{sh}$ ) are sometimes used.

The shelving-time bandwidth product indicates by how many pulse-widths a pulse can be delayed, indicating the number of operations during which a memory can hold information. A more stringent figure is given by the active Time-Bandwidth Product (TBP), which indicates how many independent pulses, or modes, a memory can store efficiently. It is then a measure of information capacity. This is the more common usage for a memory's TBP, and is useful whether we aim to store modes that are independent in time (that is a pulse train), or independent in frequency.



The exact value of  $\mathcal{TB}$  of course depends on the definitions of  $\mathcal{T}$  and  $\mathcal{B}$ , but there appears to be no uniform agreement in the literature. This is likely due in part because both quantum mechanics and optical/telecommunications engineering disciplines are involved in the study of quantum memory. As a measure of width, quantum mechanics tends to emphasise the minimum uncertainty product, and thus the variance or standard-deviation. Engineering literature on the other hand emphasises the FWHM.

In practise we're only interested in rough values of the TBP to indicate the capacity for temporal/frequency modes. In this thesis we'll mainly use the FWHM, with a notable exception of the  $1/e$  point for exponential decays (for convenience).

In addition to the TBP, the mode-capacity of a memory more generally may also take into account any other independent modes, e.g. spatial. This is particularly relevant for free-space atomic ensembles where the storage is inherently parallel in space. It's possible to store an image consisting of a large number of spatial mode 'pixels'.

## 2.2.6 Background and Reabsorption

If non-contributing background absorption  $\beta_{bg}$  is present, it will decrease the overall efficiency correspondingly, that is  $\eta_{bg} = e^{-\beta_{bg}}$ .

A very similar effect is due to reabsorption. To implement an efficient memory, the output step must mirror the absorption step. This means the output must sometimes be retrieved in the opposite direction to the input (with the exceptions of a gradient echo and Electromagnetically Induced Transparency (EIT)). If the direction is not reversed, then the output will suffer increasing reabsorption as the optical depth of the medium is increased. In some sense the coupling efficiency is also reduced, but for convenience we will absorb both contributions into  $\eta_{bg}$ . By balancing the coupling efficiency and reabsorption, the maximum efficiency that can be achieved in this case is 54%[57, 58]. This occurs for a total optical thickness of  $\beta = 2$ . Avoiding this limit by changing the retrieval direction was discussed in section 1.3.12.

## 2.2.7 Noise

Any extraneous, unwanted signal present at the output is noise. Noise is detrimental for classical storage, but can be accommodated if the signal is large enough. Thus for classical information, the worst kind of noise is that which scales with amplitude.

For quantum information the situation is opposite. The amplitude of quantum information doesn't increase with pulse size, and we're mainly interested in states with amplitudes comparable to the quantum noise level. In this case noise sources which scale with amplitude can usually be made negligible. We're instead usually concerned with independent Gaussian noise such as caused by random quantum processes. The level of noise that can be tolerated is discussed in section 2.7.

## 2.3 Linear, Coherent Memory Techniques

Four common techniques for implementing memories in ensembles are slow-light (EIT), a Two-Pulse Photon Echo (2PE), Controlled Reversible Inhomogeneous Broadening (CRIB), and an Atomic Frequency Comb (AFC). Details of each are given in the following sections. They may be summarised as in Fig. 2.2 in terms of the linear response of the medium in each case.

As the techniques are linear in terms of amplitude, they may take a general quantum state as input. They're naturally limited in the pulse shape or mode they may accommodate by their decay rate or efficient bandwidth, but the quantum state of the light in the mode may be anything. Thus a linear memory may be used equally well for continuous variable or discrete variable applications in the quantum domain.

A requirement for linearity is that inputs are of small amplitude, but this is not a serious limitation. By small we mean an average photon number much less than the number of absorbers involved. As the number of atoms in an ensemble is typically  $> 10^{20}$ , there is some leeway. This is especially true for quantum applications where  $\langle N \rangle \approx 1$ .

The efficiencies of the techniques are due to the same considerations, the echo techniques (the latter three) especially so. We'll limit ourselves to these echo techniques for our main discussion. We assume a configuration similar to that in Fig. 2.2: a Gaussian input pulse is stored in a memory with a top-hat absorption spectrum, and an exponential decay. The efficiency is estimated by eq. 2.1. With an effective optical depth of  $\beta$  and a decay time of  $\mathcal{T}$ , factors of  $\eta_c = (1 - e^{-\beta})^2$  and  $\eta_T = e^{-\frac{t}{\mathcal{T}}}$  are lost due to coupling efficiency and temporal decay over a storage time  $t$ . If the input pulse has FWHM of  $\mathcal{B}_p$  while the memory absorbs over a bandwidth  $\mathcal{B}$ , then a factor of  $\eta_B = \text{erf}(2\sqrt{\ln 2} \frac{\mathcal{B}}{\mathcal{B}_p})$  (the overlap between the Gaussian pulse spectrum and the top-hat memory. This assumes that the pulse is stored for a long time compared to its length) will be transmitted out-of-band. With these assumptions 2.1 becomes

$$\eta = (1 - e^{-\beta})^2 e^{-\frac{t}{\mathcal{T}}} \text{erf}(2\sqrt{\ln 2} \frac{\mathcal{B}}{\mathcal{B}_p}) \eta_{bg} \quad (2.2)$$

Without extra steps to reverse the echo retrieval direction, the effective re-absorption  $\eta_{bg}$  will limit the maximum efficiency to 54%, except for the case of a gradient echo (a variant of CRIB)[57, 58]. This factor may theoretically be completely removed if the retrieval direction can be reversed (section 1.3.12).

### 2.3.1 Electromagnetically Induced Transparency and Slow Light

The first method considered (and to date the most studied) for quantum storage is to directly use the reduced group velocity possible in a medium with sharp spectral transmission features. Such a feature may be created many ways, for example using spectral hole-burning[59], but the most commonly reported technique is EIT[60].

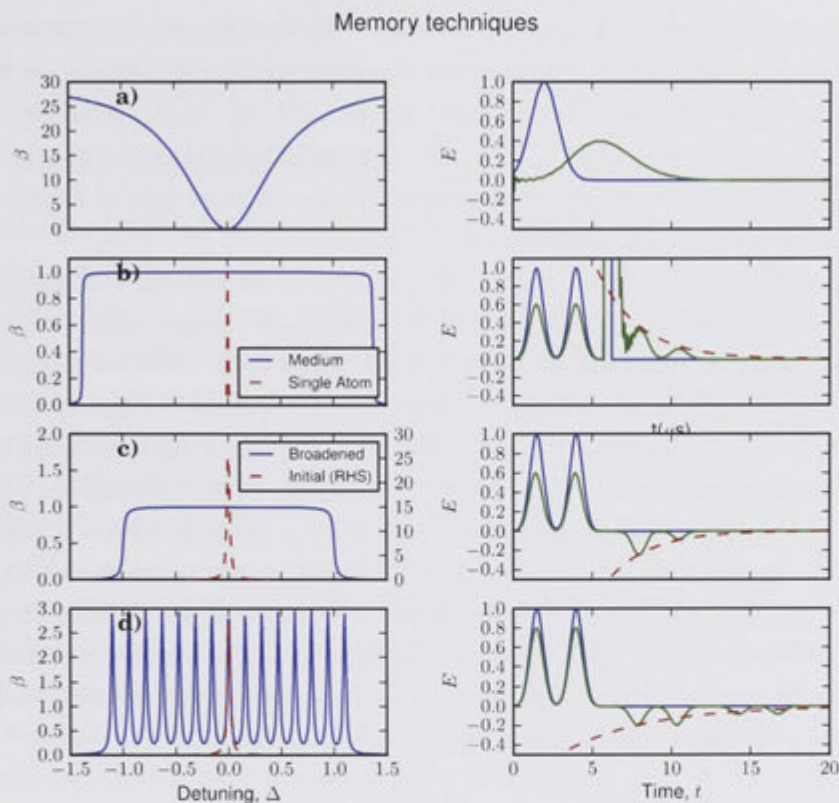


Figure 2.2: Coherent Memory Techniques. The left-hand side gives the absorption spectrum, the right hand side gives the temporal behaviour. From top to bottom the techniques are, slow-light, 2-pulse photon echo, CRIB, AFC. The latter 3 are echo techniques and have been tuned to give the same efficiency and decay time. For these, the dashed line on the left is the Fourier-transform of the temporal decay on the right. a) Slow light. A pulse is slowed when propagating through a narrow spectral hole. In an optimally efficient experiment, the hole shape is tuned in time to match the input pulse. b) 2PE. A pulse is absorbed by a broad ensemble, which is then inverted using an intense  $\pi$ -pulse. The decay time is given by the line-width of a single absorber. The echo is partly masked by an FID from the  $\pi$ -pulse c) CRIB. An initial narrow absorption feature is artificially broadened. An absorbed pulse is recalled by reversing the broadening. The decay time is given by the shape of the initial feature (although we discuss exceptions in chapter 4). d) AFC. An input pulse scatters off a comb of narrow absorption peaks, causing a rephasing at multiples of  $t = 1/\Delta$  ( $\Delta$  is the detuning between peaks). The decay time is again given by the feature width.

In EIT a narrow, transmissive spectral window (Fig. 2.2a) is created in the absorption spectrum of atoms occupying a single ground state of a  $\lambda$  level configuration (shown in Fig.2.3a). The window is created using an additional light beam, called the control, to address the empty transition. By tuning the intensity of the control beam, the width of the window may be varied. Thus the propagation of the pulse may be varied in transit. By carefully switching off the control while the pulse is within the medium, the coherence ends up between the ground states ( $|g1\rangle$  and  $|g2\rangle$ ) in Fig.2.3a). At this point the technique is referred to as “Stopped Light”. The coherence and the information may be restored by carefully ramping the control beam back on. In a sense, this type of storage is direct between the ground state hyperfine levels.

The highest reported efficiency for EIT (at the time of writing) is 42%[61], using a control beam matched to the pulse (that is the hole shape changes in time to match the pulse shape). In principle unity efficiency and arbitrarily large TBPs are possible, but so far it has mainly been experimentally considered for the storage of single modes.

The spectral feature for slow-light may be created by other methods apart from EIT and spectral hole-burning. Another example is a closely related technique to EIT, called Coherent Population Oscillation (CPO)[62]. Here a control field creates a transmissive window in a homogeneously broadened medium, and can thus be used even on phonon-broadened transitions at room temperature. The technique is not suitable for quantum applications however, for a similar reason to the standard 2PE- the medium is constantly driven between the ground and excited states so that the transmissive window is due to equal parts attenuation and amplification.

### 2.3.2 2-pulse Photon Echo

A demonstration of a 2PE[47, 63] (Fig. 2.2b) was reported only a few years after the invention of the laser. It’s known as a photon echo largely in analogy to the spin-echo[64] which at the time was already long studied in Nuclear Magnetic Resonance.

The implementation is particularly simple. An input pulse (which we assume to be small such that it’s area is  $\ll \pi/2$ ) is absorbed by an inhomogeneously broadened medium. After the initial response has died away, the medium is inverted, usually by a short intense pulse of area  $\pi$ . Each absorber receives a  $\pi$  phase shift which inverts also the sign of the phase difference between any two absorbers. As they further evolve according to their detunings, the absorbers come back into phase to radiate coherently.

The bandwidth of a 2PE is usually limited only by the bandwidth of the  $\pi$ -pulse, while the storage time is limited only by the dephasing of a single absorber as  $\mathcal{T} = \frac{T_2}{2} = \frac{1}{2\pi\Gamma}$ .  $\Gamma$  is the homogeneous line-width.

The optical depth,  $\beta$  and the decay time  $\mathcal{T}$  appearing in eq. 2.2 are just the optical depth of the medium and the dephasing rate of a single absorber.

The photon echo has been widely studied for the storage of classical information, as its only requirement is an inhomogeneously broadened medium, and it may make good use

of all the absorbers in it. Unfortunately a 2PE as it stands is not suitable for quantum storage, as the inversion causes amplification and noise. A recent discussion of this is given by [65]. However the principle of optical rephasing is not fundamentally noisy. It is possible to avoid the noise by causing the rephasing while most population in the ground state, using a total rotation of  $2\pi$  [66]. There are several ways to achieve the suppression and implement such a technique. We'll give further discussion later (section 6.2.2).

### 2.3.3 3-pulse Photon-Echo and Atomic Frequency Combs

These techniques interact the input pulse with a periodic absorption spectrum (Fig. 2.2c), to cause a fixed delay.

The delay may be understood by considering the impulse response, or FID of the peaks. Once excited by the input pulse, each begins radiating and evolving according to its detuning. The large number of them ensures they rapidly decohere, but the fixed spacing at multiples of  $\Delta$  means they will be in-phase again at times which are multiples of  $t = 1/\Delta$ . Put another way, the Fourier transform of a comb in frequency is a comb in time (similar to a spatial diffraction grating).

The earliest experiments simply used two initial pulses for preparation, the spectrum of which is a sinusoid. A third pulse can be thought of as scattering off this grating (hence 3-pulse echo). Such a spectrum may become semi-permanent in a hole-burning medium (chapter 3 discusses hole-burning).

The technique as described is simply a delay line as there is no means of changing the output time once a pulse is input. In recent years a modification has been proposed as a quantum memory by adding a shelving step, such as transferring coherence to a ground state. The time of un-shelving can be chosen, so on-demand recall is possible. Using a mechanism to reverse the recall direction is also needed to avoid reabsorption and allow high efficiencies. This more general technique, with shelving, has been called an AFC.

The decay rate,  $\gamma$  in eq. 2.2 is given by the dephasing of a single tooth. If we assume that each tooth is Lorentzian and note that the tooth separation fixes the storage time, we may immediately write the loss due to time-decay as  $\eta_T = e^{-2\pi/F}$ .  $F = \Delta/\Gamma$  is known as the comb 'finesse' ( $\Gamma$  is the line-width of a single tooth). Assuming that the bandwidth of a pulse is large enough to cover many comb 'teeth', the effective optical depth it sees during its storage time is approximately the mean absorption over the comb, that is  $\beta = \frac{\pi\beta_{pk}}{2F}$  where  $\beta_{pk}$  is the peak absorption of a single tooth.

### 2.3.4 Controlled Reversible Inhomogeneous Broadening and Optical Gradient Echoes

In CRIB the detunings of the atoms are controlled directly using an external field. The key to the technique is starting with a spectrally narrow ensemble. The external field is applied to broaden it artificially. After a pulse is incident on the broadened ensemble,

the polarity of the field is reversed. The induced detunings are then reversed, directly reversing the evolution of their relative phase.

The earliest realisations of CRIB used an electric field gradient varying with depth in the medium, and a transition with a linear stark shift to create the broadening. This is much the same as the gradient echo technique studied for decades in NMR. The difference in wavelength is significant though. In the NMR case, the spatial variation of the medium is irrelevant as the wavelength is much larger than the medium. In the optical case we are in the opposite regime, and propagation effects are very significant. As we will discuss in chapter 4, a major consequence of this is that light may be retrieved in the forward direction without suffering reabsorption. This type of CRIB technique, where there is a correlation between depth and the induced detuning, is now known as either Longitudinal CRIB or simply the name by which it is known in NMR, a gradient-echo (-memory, GEM).

Implementations of the non-depth-correlated (or “transverse”) CRIB concept, requires the use of materials where there is a continuum of different atom orientations. Examples of this are disordered materials such as glasses. When an electric field is applied, a spectrum of different shifts are induced to produce an overall broadening.

Apart from realisations using hole-burning in 2-level REI-doped solids[67, 68] and with shelving on an additional ground-state[69], a gradient-echo has been demonstrated using a feature created using an off-resonant Raman transition (as a narrow spectral hole is created using EIT, a corresponding narrow anti-hole may also be created, as in Fig. 2.3) and a magnetic field gradient[70, 71].

The bandwidth and optical depth of CRIB vary inversely as the broadening is changed. The decay time appearing in eq. 2.2 is given by the decoherence of the initial (Lorentzian) feature, giving  $\mathcal{T} = \frac{1}{2\pi\Gamma}$  where  $\Gamma$  is the line-width (we’ll discuss temporal decay with different line shapes for the gradient-echo in chapter 4). Denoting the broadening factor as  $M = \mathcal{B}/\Gamma$ , the memory optical depth is  $\beta = \frac{\pi\beta_{pk}}{2M}$  where  $\beta_{pk}$  is the absorption of the initial feature.

## 2.4 Comparison with Finite Optical Depth

The echo techniques described are essentially just different arrangements of absorption in the frequency domain. The ultimate effectiveness of each is governed by the same factors. If we allow freedom to arrange the absorption—that is the detunings of the atoms—how we like, then the maximum achievable storage time is given by the inverse of a single atoms line-width, while the bandwidth and peak efficiencies are limited by the number of atoms and the trade-off in how we arrange them. In this limit the echo techniques differ mainly in their flexibility and practicality to implement. We should also note that EIT (and other techniques based on slow-light) is essentially no different, but the requirement of temporally varying control is a significant complication to our discussion.

In an experiment with a given ensemble of atoms, the line-width of each atom is fixed,

but there may be some flexibility in how they might be arranged to maximise efficiency. There are two situations of particular interest. In memories based on spectral hole-burning (the main topic of this thesis), there is a practical limit on the maximum optical depth at any frequency for a given medium. We denote this  $\beta_{fixed}$ . A more fundamental limit is where a fixed total absorption cross-section (or a fixed total number of absorbers) is available, as will be the case if fine-grained control is possible. The total absorption cross-section may be represented by the maximum optical depth achieved if all atoms are arranged in a single spike, which we'll denote  $\beta_{spike}$ <sup>6</sup>.

For simplicity we assume a perfect experiment, including ignoring the reabsorption factor as if phase matching pulses are applied to retrieve the output in the appropriate direction for high efficiency.

If the maximum optical depth is fixed, a memory based on optical-rephasing like a 2PE will clearly be the most efficient as it may make full use of all absorbers. In a perfect experiment, there is no bandwidth limit, so we may use arbitrarily short pulses and storage times to remove the time-decay factor in eq. (2.2). This gives a peak efficiency of

$$\eta_{2PE} = (1 - e^{-\beta_{fixed}})^2 \quad (2.3)$$

To implement a memory based on an AFC we may again assume a large bandwidth. An additional loss is incurred however as some atoms must be removed to make the comb. If the remaining atoms form Lorentzian in a comb of finesse  $F$ , the efficiency is given by

$$\eta_{AFC} = (1 - e^{-\frac{\pi\beta_{fixed}}{2F}})^2 e^{-2\pi/F}$$

For a memory based on a standard CRIB technique, we can no longer assume an infinite bandwidth. The optimal configuration will be trade-off between broadening and optical depth, and the pulse-width will be of similar duration to the storage time. This case is not so simple as our assumptions in eq. (2.1) of independent time and bandwidth are no longer strictly valid, and further issues are raised about the input pulse interfering with the output. We will deal with this partially in chapter 4 (section 4.5), but in the meantime eq. (2.1) still serves as a useful approximation. If an initial Lorentzian absorption peak of width  $\Gamma$  and height  $\beta_{fixed}$  is broadened by a factor of  $M$ , the optical depth is reduced to  $\pi\beta_{fixed}/2M$  and the memory bandwidth becomes  $\mathcal{B} = M\Gamma = \frac{M}{2\pi\mathcal{T}}$ . To make reasonable use of this bandwidth and make a comparison with the other techniques we use a storage time of  $\mathcal{T}/M$ . This gives approximately

$$\eta_{CRIB} = (1 - e^{-\frac{\pi\beta_{fixed}}{2M}})^2 e^{-2\pi/M} \eta_{B,sw}$$

where  $\eta_{B,sw}$  is some loss due to out-of-band components and switching too soon, which will depend on the pulse length/bandwidth. As mentioned, simply considering the overlap of

<sup>6</sup>An equivalent limitation for  $\beta_{spike}$  would be to fix the total absorption cross-section in units of absorption depth times bandwidth

the pulse' Gaussian bandwidth shape with the memory's underestimates its bandwidth, but even with this assumption the loss is fairly minimal: a simple calculation balancing the spectral overlap (longer pulse gives better) with loss due to switching the memory before the pulse has completely entered (shorter pulse gives better) gives an optimum for the factor as  $\eta_{B,sw} = 0.8$  when the length of the pulse is  $T_p \approx 2\mathcal{T}/M$ . In practice the effective bandwidth of the memory on this time scale is larger and will increase with optical depth. The optimum pulse will then be shorter.

If we have a fixed number of atoms representing a maximum optical depth  $\beta_{spike}$ , the maximum achievable efficiency for all 3 techniques (assuming top-hat bandwidths and exponential time-decays) becomes the same, as the effective optical depth is similar so long as the atoms are arranged roughly uniformly over the pulse bandwidth. Assuming a spread of  $M$  we get the same expression as for CRIB in the maximum optical depth for a given frequency case:

$$\eta_{spike} = (1 - e^{-\frac{\pi\beta_{spike}}{2M}})^2 e^{-2\pi/M} \eta_{B,sw} \quad (2.4)$$

with the same considerations on bandwidth.

Considering the above, the choice of technique for a demonstration of an efficient quantum memory as targeted in this thesis is largely dependent on practical concerns. Barring the 2PE (which cannot be directly used as a quantum memory), the gradient-echo variant of CRIB has arguably the lowest experimental requirements. No auxiliary beams or even empty storage transitions are required. It is thus the natural choice for this work. Using Lorentzian absorption features is in general far from optimal, but it does represent a limit for the narrowest feature of most optical transitions. We'll discuss the shape of the initial feature further in chapter 4.

To achieve 50% efficiency for a simple Lorentzian case as discussed, a CRIB or AFC technique will require a medium with  $\beta \approx 20$ , or nearly 100dB absorption. In contrast, a technique based on optical rephasing could theoretically achieve this with  $\beta \approx 1$ , or 5dB. The reason for the different efficiencies may be traced to how many atoms are made use of. Eq. (2.4) indicates that the achievable TBP ( $M$ ) for a given minimum efficiency is proportional to the number of atoms available.

## 2.5 *Heralded Memories and DLCZ*

To date, implementations of quantum memories have fallen far short of unity efficiency. In a photon picture, this means the likelihood of success is low. In many contexts, the low success probability may be countered simply by trying many times- but for this to be effective it needs to be known when the storage is successful. We can't simply read the contents of a memory to check, as this will disturb it. The general problem is that of making a Quantum Non-Demolition (QND)[72, 73] measurement. When storing a single-photon, it is possible to measure its presence or not (the photon number) without destroying other properties such as polarisation, or path information. Measuring the






Figure 2.3: Schemes requiring 3 levels simultaneously. a) Electromagnetically induced transparency (section 2.3.1). The strong coupling beam ( $\Omega_c$ ) creates a transmissive window for the probe beam,  $\Omega_p$ . b) Creation of Raman absorption feature, for use in CRIB (section 2.3.4) or other techniques. A strong, off-resonant coupling beam on the auxiliary transition creates a narrow absorption where there would normally be none. In the limit of large detuning compared to Rabi frequency of the coupling beam ( $\Omega_c$ ), light at the probe frequency is effectively exciting the  $|g1\rangle \rightarrow |g2\rangle$  transition. c) Raman emission in the DLCZ protocol (section 2.5). The ensemble is coerced into radiating using a strong (usually) off-resonant pulse. A photon counter is used to record the coherent emission, collapsing it into a Fock-state. This emission is read out with another ‘read-beam’, related to a rephasing pulse. To act as a memory with an input, an incident photon causes a single excitation in the empty ground state which may be read out in the same way.

presence of the photon is known as heralding.

Heralded events are used in the DLCZ protocol[74]<sup>7</sup>. In the first instance, single heralded excitations are created by a stimulated Raman emission (Fig. 2.3c). Single-photon events are selected using an SPD. The detection of a photon heralds the storage of another, and one unit of excitation is transferred to  $|g2\rangle$ . The excitation may be converted back into photon with a further pulse Fig. 2.3b). In this protocol only a single mode can be stored at a time.

A single photon incident on the DLCZ memory is also stored, and can be recalled in the same way. In this case there is no heralding, but reading out two memories simultaneously and detecting only one photon may be used to herald entanglement between them (section 2.8.1).

A somewhat more direct heralding of storage has been demonstrated (although very inefficiently) for storage of photon polarisation, where an auxiliary emitted photon reveals the presence, but not the polarisation of the incident photon[75].

<sup>7</sup>The initials DLCZ refer not specifically to a quantum memory scheme, but an entire protocol for implementing repeaters and quantum networks using atomic ensembles. We’ll discuss this in the next section.

## 2.6 Low-finesse Cavities

A simple way of increase the coupling strength between atoms and light (i.e. the optical depth) is to make multiple passes on the same atoms. This may be done by putting the ensemble into a multi-pass cell, increasing the optical depth seen by the light many-fold.

Attention has also been given to the use of a low finesse cavity[76] to improve coupling to a single ‘pass’ of an ensemble. Large improvements to the atom-light coupling are predicted for optical depths that are matched to the cavity reflectivity.

Low-finesse cavities represent an intermediate regime— it is difficult to make high finesse cavities to couple to a single atom, and difficult to make number densities high enough to couple on a single pass. The intermediate case may be a good compromise.

## 2.7 Quantum Benchmarks and Performance Measures

Theoretical limits are one thing. In practice we wish to measure how well a given memory (or other device) is performing for quantum purposes, or whether it is performing as a quantum memory at all. The natural benchmarks for this purpose are a comparison to the classical devices we described in section 2.1: performance better than a classical memory tells us that a memory may preserve quantum information. The performance of the classical memory sets the “classical limit”. A stricter comparison is with one output of a (symmetric)  $1 \rightarrow 2$  cloning device (Fig. 2.1c). This tells us if the memory leaves behind (or otherwise loses) more information than is retrieved. The associated benchmark is known as the “no-cloning limit”.

Several measures exist to quantify how a device compares with these limits. The usefulness of a particular measure is related to its ease of measurement through experiment. Such an experiment will necessarily involve repeated storage of known states and measurements of the output, or of interference with some other reference state. A full analysis of memory performance requires the input and measurement of many repetitions of many different kinds and sizes of state. For what follows we will consider the much simplified situation where a memory behaves linearly. With this assumption we may extrapolate with repeated measurements of 1 or 2 different input types.

Common measures of memory performance are its efficiency, a level of added noise, and the quantum ‘Fidelity’ for a range of inputs. A significant complication for comparing different memories and storage of different states is the different *gains* of real devices associated with non-unit efficiency. A conceptually simple way around this problem is to introduce an imaginary pre-amplification step that normalises the gain. This gives a convenient way to define our limits for non-conditional measurements. We should note however that real pre-amplification is unlikely to be desirable in experiments as it is necessarily irreversible and will destroy some level of information. We note further that for conditional measurements such as common for single photons, even imaginary amplification is not appropriate. We’ll discuss separately the measure of fringe-visibility in this context,

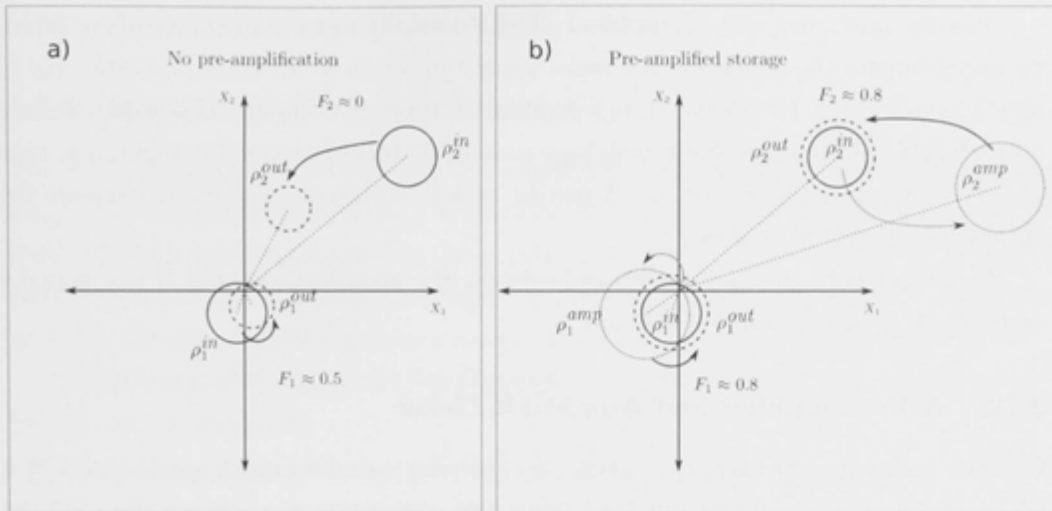


Figure 2.4: An easy way to compare memories with different gains is to imagine correcting for the loss using a quantum-limited amplifier. A memory with efficiency  $\eta$  with a pre-amplification by  $1/\eta$  has a net gain of 1. States may then easily be compared by their noise properties, or by the resultant Fidelity. The ball-and-stick plots above show the amplitude (the stick) and uncertainty (the balls) of a Gaussian state. The balls are related to the contours of a states Wigner function, such that the Fidelity between two states is related to how much their balls overlap. The states plotted correspond to an input before and after storage in a memory with an efficiency of 0.3, which also imparts a phase shift. In the unmodified case (left), the Fidelity between the input and output is nearly zero. If we first pre-amplify to account for the loss and phase shift (right), we get an overlap (Fidelity) that is independent of the input size.

which is an often (although certainly not the only) used measure in such experiments.

### 2.7.1 A Hypothetical Pre-amplifier

A direct comparison to the reference devices (section 2.1), or a direct comparison between different experiments is complicated by the different gains of real devices, caused by imperfect efficiencies. Our reference devices output pulses which look like the inputs, except for some added uncertainty, whereas a memory may not only reduce (or even increase) the size of its input, it may distort it in some other reversible manner such as phase. This is illustrated in Fig. 2.4 in phase space for a linear memory with two different sized coherent inputs. The similarity of the outputs in phase space depends strongly on the inputs.

Two natural solutions present themselves: one is to consider the ability of a classical memory to create smaller, phase shifted outputs. The approach we take is hopefully more intuitive: we imagine correcting our memory to mimic a device with unity gain, giving us an output that actually *looks* like the input. As long as the memory is linear, this may be done simply and optimally by imagining a pre-amplification step. This also is shown in phase space in Fig. 2.4. This approach is also taken by Coudreau et al. [77] for the classical limit, although it is as easily applied to the no-cloning limit as we will show.

Once we have a memory (+amplifier) which minimally distorts the mode of our input, we may compare the output to the input simply in terms of its noise properties (or in some circumstances, the Fidelity), independent of the type of input. This is particularly convenient when the input states may have a range of different coherent amplitudes, such as for continuous variable storage. A further benefit is that we may easily compare the outputs of different memories.

We should note that for many applications, the amplification step is not actually desirable. It is merely a conceptual tool.

### 2.7.2 A Pre-amplifier and Amplitude Noise

Consider an input state with  $\langle a \rangle = \frac{1}{2}(\langle X_1 \rangle + \langle X_2 \rangle)$  and variance in each quadrature  $V_{1,2}^{in} = \langle \Delta X_{1,2} \rangle^2$ . By pre-amplification and subsequent re-attenuation as above we effectively let in  $1 - \eta$  worth of vacuum twice ( $\eta$  is the memory efficiency). This is the noise penalty for inefficiency. Adding this to any noise associated with the storage,  $V^{add}$ , we have  $V_{1,2}^{out} = V_{1,2}^{in} + 2(1 - \eta) + V_{1,2}^{add}$ . For simplicity we are assuming a symmetrical treatment of the quadratures, and ignoring any phase-shift due to storage. The second two terms are the deviation from perfect storage. Comparing this to the minimum amplitude noise introduced by our reference devices of 1 and 2 for the classical memory and  $1 \rightarrow 2$  cloner, we simply get (in units of vacuum noise)

$$V_{1,2}^{add} < 2\eta \quad (2.5)$$

to beat a phase-insensitive classical memory, or

$$V_{1,2}^{add} < 2\eta - 1$$

to beat a  $1 \rightarrow 2$  phase-insensitive cloner.

In the more general phase-sensitive problem, the storage may affect the noise in each quadrature asymmetrically. To allow for this we must also generalise our amplifier and let it have a squeezed vacuum as its auxiliary input. Thus the amplifier introduces the phase sensitive  $V_{1,2}^{aux}$  where  $V_1^{aux}V_2^{aux} = 1$ . The balance is chosen to minimise the uncertainty product of the output quadrature noises. With a little fiddling, we regain the above expressions but with the more general  $\sqrt{V_1^{add}V_2^{add}}$  on the left hand side (i.e. we could have just multiplied the expressions for each quadrature).

We have derived our limits by considering how our memory may mimic our classical devices. We could alternatively have derived this the other way, by considering how the devices may try to mimic the output of our memory[78].

The above expressions may also be derived from broader considerations, motivated by criteria developed for Quantum Non-demolition experiments and teleportation. A popular example is the ‘T-V’ criteria of Ralph et al.[79]. They consider separately the preservation of signal-noise ratio,  $T$ , in combination with the noise properties of the output conditioned

on those of the input, the conditional variance  $V$ .

### 2.7.3 Quantum Fidelity for Memories

The term ‘‘Fidelity’’ (we use a capital ‘F’ to emphasise the difference from common English usage- depending on requirements it is possible to have a very good fidelity but a very low Fidelity) has a specific meaning in the context of probability theory, and (at least) two additional meanings in a quantum context. Despite the ambiguity, a quantum Fidelity is probably the most common way to measure the similarity between quantum states.

In standard probability theory, the fidelity for random variables  $x = (x_1, x_2, \dots, x_n), y = (y_1, y_2, \dots, y_n)$  is defined as

$$F_{class} = \sum_i \sqrt{x_i y_i}$$

This amounts to the correlation between the two probability distributions. If  $x = y$ , then  $F_{class} = 1$ .

The quantum mechanical Fidelity is analogous, and is defined for states  $\alpha$  and  $\beta$  as  $F(\alpha, \beta) = |\langle \alpha | \beta \rangle|^2$ , that is the transition probability. An extension including mixed states may be defined for the corresponding density matrices as [80]<sup>8</sup>

$$F(\rho, \sigma) = Tr(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})^2$$

Unfortunately another convention referring to the square-root as the Fidelity is also common in the literature.

Apart from the ambiguity of definition, the quantum Fidelity is problematic as a measure of similarity because it is very state-dependent. For instance a large, noisy thermal state may be stored with high Fidelity by a memory even if it adds further noise, or a vacuum state may be ‘stored’ with perfect Fidelity by a low efficiency memory (for instance a brick). On the other hand, a (reversible) displacement in phase, for example a phase shift between two otherwise identical states, may result in a Fidelity near zero as the overlap in phase space is now zero. This illustrated in Fig. 2.4a), where the large input will be measured as a small Fidelity, while the small input will be large.

Never-the-less, a mean Fidelity *limit* for a given distribution of inputs may be obtained. This must take in to account any prior knowledge about the type of input. For quantum communication using coherent states, it was shown that if inputs may have any amplitude, the best average Fidelity which may be achieved using a classical channel is  $F = 1/2$ [81]. For a  $1 \rightarrow 2$  cloning device it is  $2/3$  when only Gaussian operations<sup>9</sup> are allowed, or  $\approx 0.6826$ [82] for the general case. If the input ‘alphabet’ is restricted so that we can’t use a state of any size, or if squeezed/noisy/number-state inputs are allowed, these values

<sup>8</sup>The Fidelity may be defined equivalently as an overlap integral between the Wigner or Q functions representing each state in phase space (The noise ellipses in Fig. 2.4 may represent the contours of each state’s Wigner function).

<sup>9</sup>Operations that retain the Gaussian nature of the state, like amplifying, (beam-)splitting, squeezing

must be adjusted accordingly.

Our gain corrected memory agrees with these values. The Fidelity between a coherent state of amplitude  $\langle X \rangle = x_1 + ix_2$  and an output Gaussian state with  $\langle Y \rangle = y_1 + iy_2$ , and with added quadrature noises  $V_1^{add}, V_2^{add}$  is (as in Julsgaard et al.[83], noting they use a different amplitude convention of  $X = \frac{1}{\sqrt{2}}(a + a^\dagger)$ ):

$$F = 2 \exp\left(-\frac{(x_1 - y_1)^2}{2 + V_1^{add}} - \frac{(x_2 - y_2)^2}{2 + V_2^{add}}\right) / \sqrt{(2 + V_1^{add})(2 + V_2^{add})}$$

For pre-amplified states the mean amplitudes are the same, while the added noises are 2 and 1 for the classical and no-cloning limits respectively, giving  $F = 1/2$  and  $F = 2/3$  independent of the coherent state that is input (and thus the average over any set of coherent states is the same). Again however, the limits must be adjusted if noisy or otherwise different inputs are allowed.

The Fidelity more useful as a conditional measure for comparing qubit states. The classical and no-cloning limits for qubits as an average Fidelity are  $F = 2/3$  and  $F = 5/6$ [84].

### 2.7.4 Noise Corrected Efficiency

For comparing memories with different noise and efficiency characteristics, we can define a simple measure of ‘noise-corrected efficiency’. For a given noisy memory, it is the efficiency of a noise-free memory which produces the same output when pre-amplified to unity gain as discussed. We simply require the noisy and the quiet memory to have the same output variance such that  $2(1 - \eta) + V^{add} = 2(1 - \eta_{nc})$ . Thus the noise-corrected efficiency is

$$\eta_{nc} = \eta - V^{add}/2$$

( $V^{add} = \sqrt{V_1^{add}V_2^{add}}$  for the phase-sensitive case). This has some appeal as a measure with an intuitive interpretation, and gives the same quantum and no-cloning limits as other quadrature based criteria.

However we should remember that any amplification is necessarily lossy and irreversible. This means that  $\eta_{nc}$  will not be a good measure for all purposes. For example, a noiseless ( $V^{add} = 0$ ), 50% efficient device and a somewhat noisy ( $V^{add} = 1$ ), but otherwise perfectly efficient device have the same  $\eta_{nc}$  and both operate just at the no-cloning limit. By introducing an amplifier before the first device, it could be made to operate like the second, but not the other way around. Thus noise is generally more limiting than inefficiency. Some quantum properties such as entanglement may survive inefficiency, but not adding noise.

Finally, the same assumptions must hold as for other noise based criteria—the amplitude response must be linear, and the added noise must be independent of the input. Higher-order checks of this might involve measuring  $\eta_{nc}$  with different amplitudes, but the assumptions are generally very good for the small amplitude cases studied.

### 2.7.5 Conditional Measurements and Fringe-visibility

For single-photon states, our benchmarks based on quadrature noise also hold, but do not account for conditional measurements. A conditional measurement is where a result is kept (or modified) based on a conditioning event. These are especially common for single photons qubits, where the conditioning event can be whether a photon was measured at all.

A direct measure of Fidelity is to interfere the (conditionally selected) output qubits with reference states. Scanning the phase of the reference states gives interference fringes. The visibility of the fringes is defined by

$$V = (I_{max} - I_{min}) / (I_{max} + I_{min})$$

with  $I_{max}$  and  $I_{min}$  being the extrema of the observed coincidence counts. The visibility has a close relationship to the Fidelity between the states as

$$F = (1 + V) / 2$$

In many experiments, instead of single photons, the visibility is measured using coherent inputs. This is a sufficient measurement for many purposes, so long as the input state mimics the intensity of a single photon,  $\langle N \rangle = 1$ . A conditional measurement is mimicked by attenuating the reference states according to the inefficiency of the memory. We note that the temporal or spatial mode of the reference may be further modified to back-out such (theoretically reversible) imperfections in the memory. The resultant fringe pattern will then give a good measure of the memory quality subject to the same assumptions as the previous benchmarks: linearity in input amplitude, and purely independent, additive noise.

We may calculate how the additional noise in our model of a memory may affect such a measurement. If we have single photon input states, the magnitude of the observed fringes is  $2\eta$  photons per measurement. If our memory is adding an amplitude noise of  $V^{add}$ , this corresponds to a mean added photon number of  $V^{add}/2$ . This gives the level that the fringe pattern of the noisy memory is displaced from zero. The fringe visibility is then  $2\eta / (2\eta + V^{add})$  while the Fidelity is

$$F = (1 + \frac{2\eta}{2\eta + V^{add}}) / 2$$

Comparing this to the Fidelity benchmarks for qubits (2/3, 5/6), we get  $V^{add} < 4\eta$  and  $V^{add} < \eta$  for the classical and no-cloning limits when conditional measurements are allowed. Unlike for non-conditional measurements, there is no efficiency cut-off for the no-cloning limit. As long as the memory does not add noise, any photon that is measured is a useful one.

### 2.7.6 State and Process Tomography

For the more general problem of how a memory influences a quantum state, full Quantum State Tomography (QST) may be performed on the output state[85]. This is necessary if we can't assume that the memory is linear. By performing QST on the outputs corresponding to a full range of input states (a 'quorum'), a complete description of the memory may be obtained. This is known as Quantum Process Tomography (QPT)[86, 87]. Obtaining this full description is not a simple task and becomes prohibitively complex rapidly as the size of the system is increased. For smaller systems however the procedure is well studied.

## 2.8 Memories in Quantum Repeaters

As quantum communication is the main near-term practical application of quantum memory, we should try to understand the requirements for a quantum repeater. Complicating this however is that a large range of protocols and architectures have been proposed with a corresponding range of requirements. We try here to summarise only the variations relevant for memory implementation, noting that a recent review paper by Sangouard et al. [88] has much more detail. We will find that an ability to re-order modes within a memory and to read out states only partially are useful abilities, while the storage time is probably the most difficult obstacle to a practical device.

Most protocols using photons have some similarity to the DLCZ protocol (which is in turn related to the originally proposed repeater architecture[22]). We may illustrate the problem by considering the simplest communication link, divided into 2 sections. A repeater between them consists of 2 multi-mode quantum memories<sup>10</sup>.

The creation of entanglement between the end-nodes occurs in two steps: long distance entanglement is generated separately in the left and right sub-links, with each side of the repeater operating independently. The existence of measurement stations in the middle of each sub-link is usually assumed. Following this is the entanglement swapping step. Joint measurements are made between the two halves of the repeater entangles the two sub-lengths. Additional modes (or channels) usually operate independently, but there is significant benefit if they may be joined.

The main variations to the DLCZ strategy (to the best of the author's knowledge) are in the separation of entanglement generation and storage, the type of entanglement used, the way it is swapped, and the way the different modes may be combined. These are summarised in Fig. 2.5.

<sup>10</sup>The repeater is usually also split into two and the entanglement swapping done non-locally. This gives 3 links overall. This is not always beneficial though, and it complicates the problem.



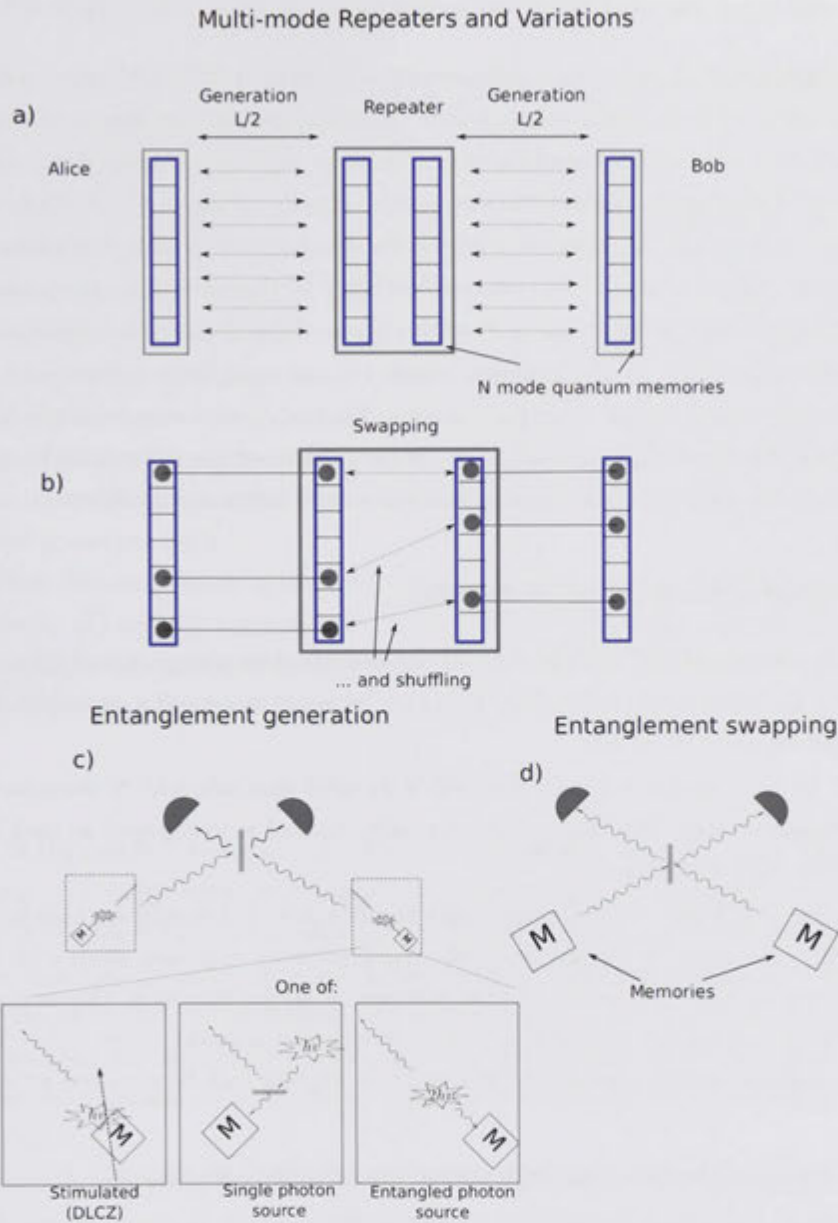


Figure 2.5: A selective summary of repeater architecture variations. a) A generic multi-mode ( $N = 7$  here) entanglement distribution architecture. We may think of the two central memories as a single repeater unit. We try to generate entanglement in all modes of each link simultaneously. Most attempts will fail (due fundamentally to loss in transmission, although there is no shortage of other loss mechanisms). b) The entanglement swapping step. If our modes are independent, we must wait until the same mode is entangled in both links. If they can be shuffled, we can swap entanglement on any modes that have succeeded. c) Different ways of entanglement generation, using the DLCZ generation or an external source. d) The entanglement swapping step is the same principle as the generation, involving a joint measurement which erases ‘which-way’ information. Doing this with single-photon entanglement requires a low success probability to avoid two-photon events.

### 2.8.1 Generation and Swapping

In the original DLCZ protocol, entanglement is generated in combination with storage, using an extra SPD to select events where single excitations are stored in the memory. These are read out simultaneously and combined at a central beam-splitter, eliminating ‘which-way’ information. Rather than using the Raman emission and SPD to create the entanglement however, an external device may also be used (which could also consist of DLCZ style photon source). This allows flexibility in the type of memory that may be used, and also allows using a source of more robust multi-photon entanglement.

Whatever the kind of entanglement used, the swapping step is basically the same: neighbouring memories are read out in some kind of joint measurement which erases information distinguishing each memory. Some post-selection might also be performed, e.g. to combine single-photon entanglement into two-photon entanglement.

### 2.8.2 Modes Shuffled or in Parallel

If our link consists only of a single channel, entanglement swapping cannot take place until both links are successfully entangled. This may be many steps and a correspondingly long time if the success rate is low.

If we have  $N$  modes, we may think of  $N$  parallel channels and  $N$  chances of success each generation step. The success rate for each channel is unchanged as each must still wait for its neighbouring link.

If we may rearrange or shuffle the modes at will however, a significant benefit is realised as was first discussed by Collins et al.[89]. Each memory may now swap with any of its neighbours that have succeeded. As the number of modes increases, this is a significant improvement, especially dramatic in terms of required memory time. If  $1/N$  is much smaller than the success probability, the need for waiting may be removed altogether.

### 2.8.3 Single-Photon Entanglement and Partial Readout

To be useful, two single-photon entangled modes must be sent for each eventual bit. This is because it is difficult to measure single-photon entanglement in an arbitrary basis as required for teleportation. Two single-photon modes are recombined at a beam splitter to achieve this.

Despite this, sharing single-photon entanglement remains attractive because of its ease of implementation. An additional issue however is that it requires, in a sense, a low efficiency process for both swapping and generation steps. Reading the memories at high efficiency erases them, and we cannot generate the entanglement at high efficiency as there is no quantum uncertainty. This is circumvented by using a low probability process repeated many times until success is recorded. The ability to partially read-out a memory many times can be beneficial for this use.

## 2.9 Meeting the Requirements

The requirements of quantum memory for repeaters were reviewed at length by Sangouard et al.[88]. Their rough consensus, on reviewing several strategies, was that efficiencies  $> 90\%$  are required for multi-mode capacities  $\sim 100$  to outperform direct strategies over distances of several hundred kilometres. This gave entanglement generation rates of 1Hz or less, and thus a correspondingly long shelving time requirement of  $> 1s$ .

The best performing protocol was found to be that of Simon et al. [90] using photon pair sources and multiple temporal modes. This does not however take advantage of the time advantage when shuffling modes. Sangouard et al. state that there is no known method to shuffle temporal modes, but this is not true. It was first demonstrated in 2009[71] using a gradient echo technique, and is indeed possible using any of the echo techniques for shuffling temporal modes as well as partial readout. This will be discussed this further in section 4.4.6.

We'll also discuss methods of increasing the number of modes in chapter 6. This allows the possibility of reaching a regime where there are enough modes that on average 1 or more entangled modes are created every step. We may express this condition very roughly as

$$N\eta e^{aL/M} > 1$$

where  $N$  is the number of modes,  $aL/M$  is the transmission loss of a single link, and  $\eta$  is the loss due to a single instance of generation and swapping.  $\eta$  is limited by the memory efficiency in the ideal case of no other losses. For a link distance of 100km and low loss otherwise, this may be as low as a few hundred modes. The relevance is that if we allow full shuffling of modes, we remove all the waiting steps. The requirement on memory time in this case may be as low as that taken for the classical heralding signals. Over a total distance of 1000km this is  $\approx 10ms$ . This represents an absolute lower limit on useful memory time.

To reach the long storage times of  $> 1s$ , the only media with demonstrated coherence and storage times of similar order long are solids[91], ion-traps[92], and more recently Bose-Einstein Condensates (BECs)[93]. To reach the absolute lower limit of 10ms, there are many more options, including optical lattices as highlighted by impressive recent demonstrations[94]. To achieve the number of efficient, shuffle-able modes for this is more difficult, although the potential for a strategy based on spatial multiplexing might offer exceptional potential. For the ensemble based strategies so far described, spatial multiplexing is possible, while the number of temporal modes is governed largely by the available optical depth. Thus for reasons of mode capacity and long-term storage, REIs in solids are good candidates for implementing quantum repeaters.

---

# Rare-earth Ions and Spectral Hole-burning for Memories

---

The anomalously sharp spectral lines of the lanthanides, or rare-earths, were first noted by Becquerel in 1908[95]. Although he was studying salts in solution, the narrow lines he saw looked like those due to a gas. It was just this ‘frozen gas’ like behaviour, combining narrow lines of gases with high densities of solids that motivated their study in the following decades. It’s these same properties that have motivated their study for storing and manipulating classical light fields, and which drive current efforts for doing the same with quantum fields. Current efforts focus on Rare-earth ion (REI)s as dopants in crystals, held at round liquid helium temperature ( $\sim 4\text{K}$ ).

With the development of the quantum theory of atoms, the narrow spectral lines were designated as  $4f^N \rightarrow 4f^N$  transitions, the 4f shell being filled as we move up the lanthanide group. The high angular momentum of the f-shell (the first to appear on the periodic table) strongly breaks the Aufbau principle. The 4fs are thus energetically higher, but physically within the already complete 5s and 5p sub-shells (corresponding to the noble gas Xenon). These protecting shells remain filled even in the triply ionised state common to the rare earths (the neutral atoms also possess 2 6s electrons which are involved in the bonding, in addition to either a 5d or 4f electron).

Transitions between the 4f levels thus occur within a tiny, very strong Faraday cage, protecting them from their environment. This is the precise requirement for ensemble based light-atom interaction. They may be packed with very high density in a solid while maintaining good coherence properties. The degree of shielding is unique to the rare earths; the element series filling the first d-shell (transition metals) and the second f-shell (actinides) display similar properties as the lanthanides, but the shielding effect is significantly weaker[96].

Further useful properties of the rare-earths stem from ground state spin-structure, particularly that due to the nucleus. In many solids, optical transitions that include a change in spin are hidden by dephasing effects. In the highly shielded rare-earths however, even weak hyperfine structure can be well resolved. This allows optical access to the nuclear spin, a perfect place for low decoherence information storage. It also allows

tailoring spectral structure using ‘hole-burning’.

However, the hyperfine levels of the REI doped solids we study are not immediately resolved. Phonon broadening is dominant at high temperature but is negligible at  $\sim 4\text{K}$ . At these low temperatures, inhomogeneous broadening remains due to static crystal disorder, much as the spectrum of a gaseous ensemble is usually much broader than a single atom due to the distribution of velocities and the Doppler shift. Similarly again to the situation in gases, we can create sharp, long-lived spectral structures within the broad line by optically manipulating the populations of the hyperfine levels. Unlike in gases, the atoms involved are not in motion, and spectral structures can persist much longer.

Manipulating the inhomogeneous line in this way is called spectral hole-burning. The method’s flexibility has allowed REIs to be considered for all manner of experiments involving the manipulation of information-carrying light. Our present focus is using hole-burning to implement the quantum memory techniques discussed in the last chapter.

The main discussion of the chapter begins with general spectroscopy (section 3.2) and the techniques and limitations of hole-burning (section 3.3). We can then address the suitability of currently available materials for hole-burning memories (section 3.4). The fundamental bandwidth limitations of these materials leads us to consider alternative techniques (section 3.5). We do this briefly as a more detailed discussion must wait until chapter 6, after we’ve discussed and confirmed relevant properties of a gradient echo technique.

To begin, we will emphasise<sup>1</sup> the potential for REIs in the broader quantum information context. After all, a memory is no use if there’s nothing you want to store in it.

Broad references on REI spectroscopy are books by Wybourne, and Huefner[97, 98] on general spectroscopy, that edited by Kaplianskii and MacFarlane [99] on hole-burning spectroscopy, and edited by Liu and Jaquier[100] for a more recent overview, although quantum applications are not emphasised.

### 3.1 Why Rare-earth Ions

It is easy to motivate the use of REIs for a quantum memory demonstration, due to the flexibility of spectral hole-burning, and the long history of classical information processing in REI doped solids. A stronger motivation is that these materials are a realistic candidate for practical devices in quantum information at large.

For the creation of real, useful devices, the ideal implementation for many technologies is in a solid. This is for the practical engineering requirements of manufacturing, for robustness, self-containment and viability for miniaturisation. Such practicality will likely apply to the quantum information devices of the future. Beyond the engineering perspective, the solid state represents the extremes of high density, long-term stability and

---

<sup>1</sup>advertise!

potential coupling between quantum systems. These are requirements for strong atom-light coupling, long term quantum storage and scalable quantum computing.

Realisation of quantum processing in actual solids however is complicated by these same extremes, where the range of strong interactions means only slight level of disorder will wash out quantum effects. While strong interactions between quantum systems is needed for computing purposes, they must be ordered and controllable. In most solids, this is impossible to achieve.

Two directions toward taming these interactions may be identified: removing them or reducing sensitivity to them. The first is emphasised by approaches such as phosphor-in-silicon quantum processing[101]. This is made possible by the properties of silicon and the technological manufacturing achievements of the semi-conductor industry. The purity and order of these crystals is unmatched elsewhere. In such a pure medium, quantum systems (the phosphors) are doped back in an ordered, controllable manner. This approach has many benefits of control and use of existing technology, but does suffer in the sense that it cannot realise the full density of systems possible in a solid.

The use of REIs in solids takes the second approach, avoiding the sensitivity to disorder. These ions may well have the best shielded electron transitions of any solid-state optical centre. Due to the strength of shielding, they may be packed close together in a wider range of solids, while maintaining isolation at the quantum level. They are also much less sensitive to material preparation conditions. An approach to quantum computing using such systems might involve selecting sub-groups which have naturally occurred in the right configuration, rather than creating them explicitly. Other benefits of REIs are a relatively good nuclear isolation and good chemical flexibility.

Of course it may be ideal to combine the approaches, using well shielded transitions in low disorder, well-controlled systems. Perhaps the main issues are the poor size and chemistry match for REIs in silicon, and the lack of technology and industry for growing low-disorder REI crystals. This remains a very promising direction which will merit later attention (sections 3.5.3, 6.2.4).

## 3.2 Spectra of 4f-4f Transitions in Crystals

This is a brief, non-mathematical overview of 4f levels and their spectroscopy. The focus is on transitions with high coherence and long-lifetime as suitable for quantum memory. A much more general and in-depth introductions is given by Liu[100].

The many lanthanide species and the various types of crystals they can be doped into means a large range of spectra and interactions that might be considered. Our range of interest is much narrower as most information processing experiments consider a similar regime: where the shielding is good with respect to environmental perturbations. This is in contrast to many industrial uses of rare earths where rapid, incoherent interactions are emphasised.

We first describe the spectroscopic picture seen in most information processing experiments, before briefly considering the physics leading to this picture. We'll then discuss the main parameters of experimental interest and their variation.

### 3.2.1 A Typical Spectrum of a Rare-earth Doped Crystal

REIs are used in industry for a wide array of applications. The properties of the 4f levels are important for many of these, but the conditions are usually very different from our situation.

For information processing applications, it is mainly certain optical transitions of  $\text{Pr}^{3+}$ ,  $\text{Nd}^{3+}$ ,  $\text{Eu}^{3+}$ ,  $\text{Tm}^{3+}$  and  $\text{Er}^{3+}$  that have been targeted. These elements are doped into good quality, magnetically quiet crystals at the  $<1\%$  level and held at cryogenic temperatures ( $\sim 4\text{K}$ ). The narrowest, most coherent optical transition (or infra-red for  $\text{Tm}^{3+}$ ,  $\text{Er}^{3+}$ ) is that studied.

The spectrum in Fig. 3.1 illustrates a common situation for the studied materials. It is biased toward our present work, but it applies fairly generally apart from the number of ground state levels and overall scales of broadening. Fig. 3.2 illustrates temporal dynamics.

The spectrum of a single dopant ion consists of narrow lines similar to gaseous atoms, the number of lines corresponding to hyperfine (or sometimes Zeeman) ground states at the level of  $1\text{MHz}$ – $1\text{GHz}$ . At equilibrium and at standard cryogenic temperatures, the ground states are equally populated. This is simply because the available thermal energy ( $\kappa T$ ) is comparable to the transition energies transitions only when  $T \sim 1\text{mK}$ – $1\text{K}$ .

The optical transition of each ion in the crystal occurs at a slightly different detuning due to variations in its local environment. This causes the spectrum of the crystal as a whole to appear as a single, much broader line (Fig. 3.1, top).

### 3.2.2 Free-ion Approximation

The strong shielding of the  $4f^N \rightarrow 4f^N$  transitions greatly simplifies the spectroscopy of REIs doped into solids as the effect of the host is usually small enough to be treated as a perturbation. The ion is modelled as sitting in a potential well due to the host, called the "crystal-field", which perturbs the state of the otherwise free ion. This fact has been a strong motivation for studying REIs, as it has allowed the study of interactions in the solid-state that are otherwise too complex to decipher.

The rough magnitudes of the interaction energies determining the 4f energy levels in a solid were tabulated by Liu[100]:




Figure 3.1: A typical absorption spectrum and important parameters of a REI doped crystal held at 4K. Most materials studied in QI applications are qualitatively similar. Left: the level structure is usually split according to spin and external magnetic field (here into 3 levels, no external field). At equilibrium and normal temperatures, each of the ground states is equally occupied. Right top: The spectrum of the crystal appears as a broad, Gaussian line on the order of GHz wide. This is made up of the narrow but randomly shifted spectra of single ions. Right lower: The spectrum of a single ion, coloured according to which ground state it is in.



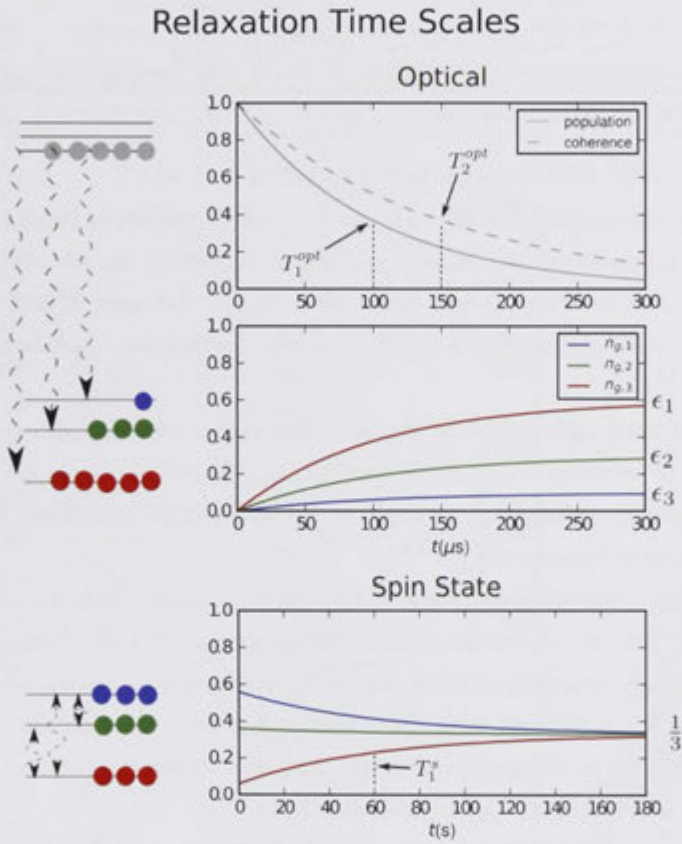


Figure 3.2: Relaxation time-scales. We start with a population of ions in a single excited state, from which they spontaneously decay with a time constant  $T_1$ . The branching ratios  $\epsilon_n$  from the particular excited state govern the probability of decay into each ground state. On a much longer time-scale, the ground states will re-thermalise until each level is equally filled. The lifetime of the states sets an upper limit on the coherence time between them of  $T_2 < 2T_1$ . For optical transitions this limit is often approached, at which point they are said to be ‘lifetime limited’. For hyperfine levels, coherence times remain much shorter than the (very long) lifetime limit unless decoherence controlling techniques are used (section 3.2.10).

Interaction	Energy (cm <sup>-1</sup> )	Frequency
Configuration splitting (transition to 5d level)	10 <sup>5</sup>	10 <sup>15</sup> Hz
Intra Configuration Splitting (4f-4f)		
Non-central Electrostatic field	10 <sup>4</sup>	100THz
Spin-orbit	10 <sup>3</sup>	10THz
Crystal field	10 <sup>2</sup>	1THz
Ion-ion induced band structure (stoichiometric only)	10 <sup>-2</sup> → 10	100MHz → 10GHz
Hyperfine interaction(nuclear spin)	10 <sup>-3</sup> → 10 <sup>-1</sup>	10MHz → 1GHz
Super-hyperfine interaction (neighbouring spin)	10 <sup>-4</sup> → 10 <sup>-2</sup>	1MHz → 100 MHz

That is, the crystal field is only a perturbation of around 1%.

In a Coulombic potential, the 4f<sup>N</sup> orbitals would be completely degenerate (14!/(14 - N)!N! -fold). In a populated ion, these levels are split first by the non-central field of the other electrons. The splitting is into states of definite total spin  $S$  and orbital angular momentum  $L$ , each remaining  $(2S + 1)(2L + 1)$ -fold degenerate. Each separate level may be labelled as  $^{2S+1}L$ .

Each  $^{2S+1}L$  is then split again by the so-called spin-orbit interaction, essentially the relative orientation between  $S$  and  $L$ . In the free-ion, these final levels are fully described by the total angular momentum  $J$ , and are  $2J + 1$  fold degenerate. They may be labelled by a Russell-Saunders term symbol  $^{2S+1}L_J$ .

This level system describes quite accurately the large scale levels even when the ion is present in a solid. Fig. 3.3 shows the classic Dike diagram[102] published in 1968 by Dieke and co-workers. It gives measured levels for all REIs doped into LaCl<sub>3</sub>, labelled according to term symbols. A comparison with this diagram has been successfully used to identify unknown spectral lines in all manner of hosts. The host dependent crystal field only gives splitting comparable to the width of the drawn lines.

The Dieke diagram has since been extended to higher energies, and with modern computers it's been possible to theorise about the entire 4f series. An essentially complete energy level diagram was calculated by Peijzel et al. , also shown. Agreement with experiment is good, and even many of the higher energy levels have been verified using synchrotron radiation[103].

### 3.2.3 Crystal Field Levels

In the non-spherical potential of the crystal field,  $J$  is no longer a good quantum number. The  $2J + 1$ -fold degeneracy of each  $J$  level is lifted to a degree that depends on the symmetry of the dopant site. The maximum lifting occurs for a site of axial symmetry or less.

For ions with no un-paired electron spin (i.e. those with an even number of electrons, known as non-Kramers ions), the maximum lifting corresponds to a complete separation into singlets. For higher symmetry sites they may remain doublets (non-Kramer's

**Classic 'Dieke Diagram' for Rare Earth Ions**

Extension by Peijzel et. al.





Figure 3.3: (left) The original Dieke diagram with measured energy levels in  $\text{LaCl}_3$ , and (right) theoretical extension by Peijzel et al.[103]. Many of the predicted levels have been confirmed, including at higher energies using Synchrotron radiation.

doublets). For those with an odd-number of electrons (Kramers ions), the electron spin degeneracy remains and  $J + 1/2$  doublets are created.

The distinction between ions with singlet ground-states (non-Kramers ions in low symmetry sites) and those without is important as the former have no magnetic moment to first order. This makes them much less sensitive to magnetic fields.

### 3.2.4 Hyperfine Structure in Zero Field

In the absence of an external magnetic field, there is a final level of splitting for isotopes with a nuclear spin. These splittings occur also in the free-ion, but are modified by the crystal field. The strongest hyperfine splitting occurs for ions with an electronic magnetic moment, that is Kramer's ions and non-Kramer's doublets. The interactions are of order 1GHz.

For nuclei with spin  $> 1/2$  a nuclear electric quadrupole interaction is also present, causing a splitting even for non-Kramer's ions in low symmetry sites. For these ions the splittings are much smaller (1MHz  $\rightarrow$  100MHz). There are several factors determining the splitting including interactions with other f-electrons, the gradient of the crystal field at the nucleus, modification of the closed electron shells by the crystal field, and the "pseudo-quadrupole" effect resulting from a second order magnetic dipole interaction[100]. The rough degree of the splitting appears to be determined by the ion properties but with large host-dependent variation.

Finally structure may also be induced due to magnetic moments of nearby host ions. The effect of nearby nuclear spins is known as the super-hyperfine interaction. The field due to neighbouring spins is  $\sim 1$  gauss, inducing splittings of  $\sim 1$ MHz on Kramers ions or  $\sim 1$ kHz on electronic singlets with a nuclear spin.

### 3.2.5 Zeeman Structure

An external field affects any levels with a magnetic moment, the strongest of which are again electron spins and non-Kramer's doublets. If there is no nuclear spin, the external field is required to lift the associated degeneracy. Electron spin moments are of the order of the Bohr magneton ( $\sim$ MHz/G).

Further down the energy scale is the nuclear-Zeeman splitting. For ions with no electronic moment, the nuclear levels are typically much less sensitive although enhancement/desensitising factors can make a large difference to the bare nuclear moment. Two examples are quenching of the Eu nuclear moment by mixing of an electronic level, measured at  $\approx 0.3$ kHz/G YAlO<sub>3</sub> (YALO), and anti-shielding by mixing of very close crystal field level in Ho<sup>3+</sup>:LaF<sub>3</sub> with 0.25MHz/G[104].

### 3.2.6 Lifetimes

The lifetimes of the optical levels (i.e. those determined by  $L$  and  $S$ ) are limited by phonon (non-radiative) or photon (radiative) decay. The lifetime of the spin states are usually only limited by phonon decay at higher temperatures or magnetic fields, but at all temperatures and fields by cross-relaxation with other ions.

Starting with the optical (the levels determined by  $L$  and  $S$ ), most levels undergo rapid non-radiative decay due to phonons bridging the gap to the next lowest level. This remains the case even at low temperatures. The degree is reduced with softer materials, but remains the rule. The exceptions are the levels with a large energy gap (more than a few phonons worth) beneath them, in which case phonon decay may be almost completely eliminated. This is usually the case for the lowest levels of a  $J$  multiplet. In good quality crystals the lifetimes of these levels are strongly radiative, and thus determined by oscillator strength.

For the spin-levels of an optical ground state, both phonon assisted decay (known as spin-lattice relaxation) and resonant cross-relaxations with other ions play a role. Phonon decay was described by Orbach[105] as caused by a one-phonon (direct) process and a two-phonon (indirect) process. The direct process is both field and temperature dependent with  $T_1 \propto H^{-2}T^{-1}$  ( $H^{-4}T^{-1}$  for Kramers), while the indirect process depends only on temperature as  $\propto T^{-7}$  ( $T^{-9}$  for Kramers). At cryogenic temperatures and low fields, both may be very low, but become relevant very quickly as temperature or field are increased.

The resonant exchange of spin state between nearby ions are not eliminated in low fields and temperatures. Usually the only resonant spins are those belonging to ions of the same species, and so the mechanism is strongly dependent on doping concentration. Host spins may play a role mediating the interaction, but don't cross-relax directly.

In the low doping limit and at low temperatures and fields, the electron spins theoretically have much longer lifetimes than normal experiment time scales. In practise, lower times of  $\sim 10\text{ms}$ [106, 107] are measured for rare-earths.

Nuclear spins interact more weakly, and have correspondingly longer lifetimes. Nuclear Zeeman levels have been reported with lifetimes of 30s or longer[108, 109]. By far the longest lifetimes however have been demonstrated with nuclear quadrupole levels, with measurements showing minutes to hours and predictions as long as weeks[110]. The improvement with respect to nuclear Zeeman levels may be attributed to local variations in the quadrupole splitting (i.e. inhomogeneous broadening) due to variations in the crystal field. These deviations shift nearby nuclei out of resonance by several line-widths, reducing direct cross-relaxations.

### 3.2.7 Transition Strength

$4f^N \rightarrow 4f^N$  transitions are electric dipole forbidden for a free ion, occurring only weakly via the magnetic dipole interaction. Yet their observed intensities in solids are usually much greater than predicted from this alone. The cause remained unclear until the 1960s,

when a solution was reported simultaneously by Judd and Ofelt[111, 112]. The so-called Judd-Ofelt theory has since been successfully used to model most  $4f^N \rightarrow 4f^N$  transitions. An accessible introduction to the theory and its importance is given by Wybourne[113].

The lowest-lowest transitions that we are interested in occur usually by ‘forced’ electric dipole, and occasionally by magnetic dipole. In the free-ion the  $4f^N \rightarrow 4f^N$  are electric-dipole forbidden as the levels have the same parity. The asymmetry introduced by the crystal field in a solid mixes in odd-parity configurations of excited states, typically those corresponding to the 5d level. For this reason the type of host environment, and especially its symmetry, plays a large role in determining transition strength. A large transition dipole moment is associated with a large static dipole moment induced by the crystal field, i.e. Stark shift.

In general, transitions with a forced electric-dipole character occur between levels with  $\Delta J \leq 6$ , while magnetic dipole transitions also occur where  $J \leq 1$ , such as the  ${}^4I_{15/2} \rightarrow {}^4I_{13/2}$  transition of Erbium. An exception is transitions between  $J = 0$  states as for the  $\text{Eu}^{3+}$  transition commonly targeted in Europium, which remain magnetic dipole forbidden due to the lack of a magnetic moment.

### 3.2.8 Homogeneous Broadening

The observed line-width of a single ion transition is governed by its coherence time,  $\Gamma_h = 1/(\pi T_2)$ . The lifetime gives an upper limit on the coherence time,  $T_2 < 2T_1$ . An interaction which affects the transition energy will reduce  $T_2$  without reducing  $T_1$ .

The lifetime limit may be approached for optical transitions in high quality materials, but rarely for spin transitions. Dephasing of the spins results in similar dephasing on the optical, but for the optical transition this may be negligible compared to the lifetime limit. Spin states only approach the lifetime limit if the lifetime is especially short.

For lowest-lowest optical transitions whose lifetime is not influenced by phonon decay may still experience dephasing due to phonon interactions in the ground state. This occurs if the crystal field has not strongly separated the ground state from the next lowest level as a lattice vibration can induce mixing. At the other extreme, glassy materials have also received a significant level of study. These show strong dephasing over a wide range of time-scales which continue even at very low temperatures. This is modelled as re-arrangements of so-called “two-level systems”[114, 115].

Aside from lattice interactions, the main remaining dephasing mechanism may be the same for spin and optical states, due to magnetic fluctuations. Nuclear (or other) spins in the medium may flip on experimental time-scales causing a varying magnetic field seen by the dopant. The details of such interactions can be very complex as they may depend on the propagation dynamics of host spins. These can lead to complex, non-exponential decays.

The primary example of propagation dynamics is known as the frozen-core effect, which occurs when the dopant has a larger magnetic moment than ions of the host[116, 117].

The field due to the dopant detunes the nearby host spins from the larger environment, causing them to “freeze”. The core thaws slowly from the outside (least detuned) inwards, so that the dephasing rate increases with time. This also results in spectral diffusion as the detuning of the dopant spin changes with time. Spectral diffusion highlights an occasional ambiguity with definition between lifetime and coherence—it is usually thought of as a single decohering transition, but it may also be considered as a lifetime effect due to transitions between super-hyperfine levels.

The ultimate sensitivity of the dopant to magnetic dephasing is determined by its own magnetic character. Thus singlet states, such as those of non-Kramers ions in low symmetry sites are the most robust.

A final significant broadening mechanism with severe implications for classical information storage is Instantaneous Spectral Diffusion (ISD)[99]. The optically excited ion may have a different electric dipole moment, causing dephasing on those nearby. The dephasing interaction scales with the strength of atom-light interaction, as well as with the number of ions that are excited, severely limiting the information density when storing intense ‘classical’ pulses. ISD is an issue only when a significant population is excited, whereas for light intensities at the quantum level it is completely negligible. It will only become an issue as we approach a number of qubits (photons) comparable to the number of atoms in the material. In other situations, ISD is an example of a controllable, coherent interaction between ions[118].

### 3.2.9 Inhomogeneous Broadening

Having discussed the properties of individual ions, we now turn to the properties of the ensemble. While the (optical) decoherence time of an individual ion is long and hence its line-width narrow, the ensemble as a whole may be broader by a factor  $10^7$  or even more (Fig. 3.1). The inhomogeneous line (width  $\Gamma_{ih}^{opt}$ ) is due to the sum of many narrow lines at different detunings. The variations are static and thus do not cause decoherence.

The dominant physical mechanism for the broadening is variations in the crystal field. The variations are caused by any kind of disorder or defect: vacancies, fault planes, chemical (or even isotopic) impurities and the dopant ions themselves cause local variations as well as longer range strain fields. The actual level of broadening varies greatly with material and growth conditions, but variations are typically on the order of 1% of the total crystal field. This corresponds to inhomogeneous lines of  $\Gamma_{ih}^{opt} = 1\text{GHz}-100\text{GHz}$ . In high quality crystals, the major source is often the dopants themselves due to the variation in ion size.

The Zeeman spin levels are unaffected by the crystal field and thus are not broadened, but nuclear quadrupole levels see a similar relative shift to the optical. The absolute shift however is much smaller of  $\Gamma_{ih}^s \sim 10\text{kHz}$ .

### 3.2.10 Reducing Dephasing

It is often possible to reduce factors which contribute to dephasing. This is especially true for dephasing due to spin flips.

The easiest, and most common method to reduce sensitivity to flipping host spins is to apply a magnetic field. This may improve the coherence of both optical and spin levels. The field strength is comparable to the ion-host interaction, the interaction energy will change less as a result of spin-flips unless the flip is exactly aligned with the applied field. The field direction may be chosen to maximise this effect. The required strength depends on the interaction strength. It may be as low as a few gauss for non-Kramers ions, while larger fields of  $\sim 1\text{kG}$  are required for those with an electron spin.

It is possible to reduce dephasing of an optical transition due to coupling between spins of the same species (RE-RE coupling) by reducing the population of the other spin levels. While it should be possible to do this through hole-burning, a more effective (and stable) solution is possible for Kramers ions with a large magnetic moment. By apply a large field and reduce the temperature, a splitting larger than  $kT$  may be obtained so that the population thermally relaxes to the ground state[106]. This is clearly more feasible for ions like  $\text{Dy}^{3+}$  or  $\text{Er}^{3+}$  which have particularly large  $g$  factors.

Dynamic Decoherence Control (DDC)[119] describes a class of techniques for driving either the system of interest or its environment on a faster time-scale than the decoherence occurs. In the simplest case the interaction is averaged out. A wide range of techniques have been studied for improving spin coherences over many years of NMR spectroscopy. These techniques have also been applied to optical transitions in rare-earths[120, 121].

A final technique which has resulted in the recent demonstration of the longest coherence times so far reported for a REI in a solid ( $> 2\text{s}$ ) by Fraval et al. [122]. The technique is again to apply a magnetic field, but this time specifically tuned to the specific transition targeted. The transition is tuned to a turning point for changes to the field in any direction, a *critical point*. The transition then has zero magnetic sensitivity to first order, so the technique is known as Zero First-Order Zeeman (ZFOZ). Remaining dephasing is due to the second-order moment which is dependent of the flatness of the critical point used. The technique may only be applied to transitions which go through a turning point. In the recent demonstration[122] this was due to an anti-crossing between nuclear quadrupole levels. It is not applicable to pure Zeeman levels.

## 3.3 Tailored Spectra using Metastable Hole-burning

Now that we've described the spectrum, we'll discuss changing it. Spectral hole-burning is a way of making nearly arbitrary spectral structures, by manipulating the state of atoms to change their resonant frequency. There are many possible mechanisms for hole-burning (see e.g. MacFarlane [109]), but the most flexible are given by the metastable spin levels already discussed. In our discussion we have in mind hyperfine levels, particularly



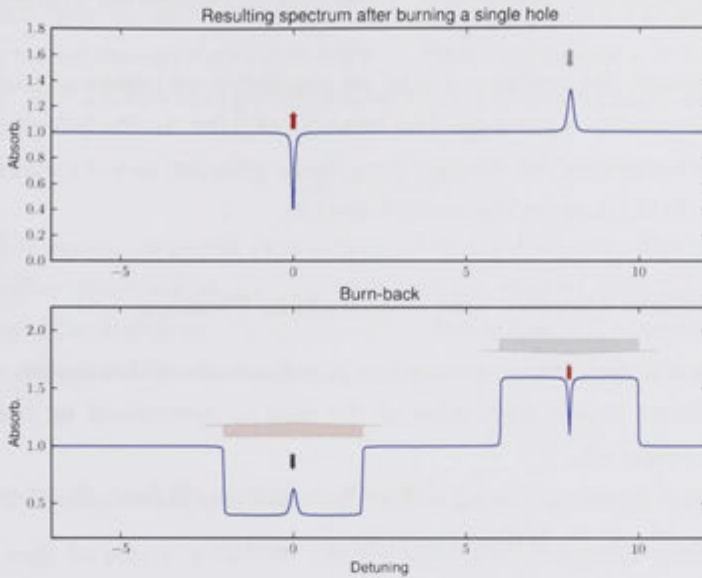


Figure 3.4: Spectral hole-burning, on a narrow spectral region of the inhomogeneous line in Fig. 3.1. Top: the laser excites a narrow packet of ions which fall into a different ground state, leaving a transmissive hole. Bottom: population shuffling to make a narrow feature, often called ‘burn-back’. An empty trench was made first, then some ions were moved back into it. This resultant feature may be slightly wider than the hole made due to inhomogeneous broadening of the ground states.

nuclear quadrupole levels, as the most effective and best studied population reservoirs. The principles remain the same for other levels, so long as the spectrum is described as in Fig. 3.1. Other mechanisms will be touched upon in section 3.5.1.

The principle of hole-burning is simple. When a spectrally narrow laser is incident on our ensemble, a spectrally narrow packet of atoms is excited. When these atoms decay, there is chance that they’ll fall into a different ground state than they started in, and thus out of resonance with the laser. If they fall back into the same ground state, they’ll be excited again. A tailored spectrum may then be created by spectral ‘carving’, simply using the laser to excite spectral regions that we wish to be transparent.

More complicated sequences may be used to create higher, sharper spectral features, to create features consisting of a certain transition or crystal site, or to get around limitations such as the inhomogeneous broadening of the spin transitions or propagation effects. The spectrum that can be made is quite flexible, but ultimately limited in sharpness by the homogeneous line-width, and in total bandwidth by the hyperfine structure.

A situation for good hole-burning may be summarised by 3 conditions: the transition is mainly inhomogeneously broadened,  $\Gamma_{ih}^{opt} \gg \Gamma_h^{opt}$ , the metastable hole-burning level is separated by many homogeneous line-widths  $\Delta_s \gg \Gamma_h^{opt}$ , and the hole-burning lifetime is

much greater than the time taken to pump ions there  $T_1^s \gg T_1^{opt}/\epsilon$  ( $\epsilon$  gives the branching ratio back to the original ground state). These may be understood with reference to Figs. 3.1 and 3.2.

We'll first discuss the specifics of how an population of inhomogeneously broadened absorbers is addressed by a given spectral intensity of light. In the following sections we'll consider specific techniques for manipulating the population and it's absorption. Finally we'll summarise limitations on the created spectra.

### 3.3.1 Population Density and Absorption Profiles

The precise way the laser addresses the ion population for hole-burning, and the way a given ion population causes absorption of the laser is determined by the homogeneous broadening of a single ion.

The absorption spectrum measured for a spectral population of ions  $P(\omega)$ , each with an absorption profile,  $L(\omega)$ , is determined by convolution:

$$A(\omega) = \int_{-\infty}^{\infty} P(\omega_0)L(\omega - \omega_0)d\omega_0$$

The absorption spectrum of a single ion is nearly always governed by a decay rate, giving a Lorentzian in frequency space  $-\Re(\kappa(\Delta)) = L(\Delta) \propto \frac{1}{\Delta^2 + (\Gamma_h^{opt}/2)^2}$  (section 1.3.2).

If  $\Gamma_h^{opt}$  is much smaller than the frequency range of interest,  $P(\omega)$  may be a sufficient indication of the line-shape.

A similar problem is to calculate which ions are excited, and thus hole-burned away in an experiment. If we consider slow burning rates, so that ions decay faster than they are excited, the interaction is incoherent. The rate of burning ( $R$ ) is dependent on the rate that energy is delivered at a given frequency, and the strength of coupling between light and atom (the Rabi-frequency squared):

$$R(\omega) = \int_{-\infty}^{\infty} I(\omega_0)L(\omega - \omega_0)d\omega_0$$

The population density  $N(\omega)$  and the laser intensity  $I(\omega)$  may take many forms, but we'll often idealised them as square edged profiles. It's convenient to remember that convolving a sharp edge with a Lorentzian smears it into an arctan shape. Specifically for the case of a top-hat profile with corners at detunings  $\Delta_1$  and  $\Delta_2$ , the burn rate is described by a spectral profile

$$H(\omega - \Delta_1) - H(\omega - \Delta_2) \implies \frac{1}{\pi} \left( \tan^{-1}\left(\frac{\omega - \Delta_1}{\Gamma}\right) - \tan^{-1}\left(\frac{\omega - \Delta_2}{\Gamma}\right) \right)$$

where  $H$  is the Heaviside step function.

The given shape describes the excitation rate due to a top-hat intensity, or the absorption due to a top-hat population, where the homogeneous line-width is  $\Gamma$ .

### 3.3.2 Power Broadening

The above holds only for relatively weak intensities of light. An exception occurs if our driving field is strong enough that atoms undergo Rabi-oscillations at a comparable rate to the line-width. The effective line-width is modified by the driving rate (the Rabi-frequency eq. (1.44)):

$$\Gamma_{\text{eff}} = \Gamma + 2\Omega$$

This effect is known as power broadening. It is a saturation effect as increasing the laser power further drives the excitation down from the excited state before they've had a chance to decay. The maximum pumping rate is already reached when the excitation rate is equal to the decay rate. Increasing the laser intensity further just broadens the spectral width effected.

### 3.3.3 Burn-back and Population Shuffling

The simplest way to create a narrow feature is to start with the inhomogeneous line, then create 2 transmissive windows with a narrow gap, leaving the absorptive feature. An alternative method is to start with a transmissive window and move absorption into it. To do this we follow the reverse principle: we find ions which *could* be absorbing where we want, and keep exciting them until they fall into the ground state where they do. We'll call this 'burn-back'. Burn-back and more sophisticated population shuffling set the laser frequency to linear combinations of the hyperfine detunings, addressing specific transitions.

The creation of a narrow feature using burn-back is shown in Fig. 3.4. We create a broad, empty spectral trench by scanning the laser. A narrow hole is then burned at a frequency offset from the centre of the trench corresponding to a hyperfine splitting. Some of the absorption from the hole will appear as an anti-hole in the centre of our trench. Usefully, this feature will consist only of ions in the ground state chosen. Sequences with more steps, but using the same principles, may further select for the excited state.

The feature created this way is generally broader and shallower than the corresponding burned hole. This is due to the inhomogeneous broadening in the hyperfine levels,  $\Gamma_{ih}^s$ , and to decay into other ground states.

### 3.3.4 Coherent Optical Transients for Hole-burning

Coherent Optical Transient (COT)s refers to coherent pulses shorter than the optical decay time of the atoms addressed (a quantum memory is an example of a COT). Hole-burning with COTs has the potential to realise more rapid, efficient pumping as they may avoid power broadening and propagate differently.

One example is a sequences of 'zero-area' pulses for narrowing an initially broad feature. They work by having an effective pulse area of zero for the atoms that are wanted in the feature, while those not wanted see a finite pulse area. Thus the wanted atoms are returned

to the ground state and are unaffected, while those unwanted remain excited after the pulse is passed and are pumped away. This has been demonstrated using pairs of a short and a long sinc-pulse for creating a top-hat feature[123, 124]. They are each scaled to have the same peak height in frequency space, but with a  $\pi$  phase difference. Thus the short pulse excites a broad spectrum, while the long one returns the wanted ions near the centre to the ground state.

In some cases COTs may be fast and efficient in comparison to incoherent burning as they avoid power broadening effects, but they are ultimately limited by the same homogeneous line-width as incoherent burning. This is because greater frequency resolution requires longer pulses, which are of course limited again by the homogeneous lifetime.

### 3.3.5 Competing Pumping

It is sometimes useful to weakly drive two transitions at a time, in particular combining the competing processes of burning back and burning away. The simplest example is applying laser power on two transitions with a common excited state. As long as the driving is weak (and preferably sequential to avoid more complicated mixing) this may be treated with rate equations. We'll use such a technique in chapter 5, partly for the purpose of removing the effect of inhomogeneous broadening on the ground states,  $\Gamma_{ih}^s$ .

We'll treat only a simple  $\lambda$  system, although the extension to more ground states is straightforward. We apply incoherent burning with a power spectrum of  $I_1(\omega)$  to the first and  $I_2(\omega)$  to the second transition. The transitions have dipole moments  $\mu_1$  and  $\mu_2$ .

Starting with only the first driving field, the rate at which ions in the first ground state are driven to the excited state is  $\Omega_1^2 = I_1(\omega)|\mu_1|^2$ . The fraction of these that decay to state 2 depends on the branching ratio,  $\epsilon$ . In the absence of other decay mechanisms, this is given by transition strengths  $\epsilon_1 = |\mu_1|^2/(|\mu_1|^2 + |\mu_2|^2)$  and  $\epsilon_2 = |\mu_2|^2/(|\mu_1|^2 + |\mu_2|^2)$ . The populations  $P_{1,2}$  in each ground state decrease due to  $I_{1,2}$  at the rate

$$R_{1,2}(\omega) = I_{1,2}(\omega)P_{1,2}(\omega) \frac{|\mu_1|^2|\mu_2|^2}{|\mu_1|^2 + |\mu_2|^2}$$

Under the influence of both processes, the spectrum will eventually stabilise when  $R_1 = R_2$ . Adding the restriction that the population sum to 1 ( $P_1 + P_2 = 1$ ), the final population in state 1 is then calculated easily according to the spectral intensity incident:

$$P_1(\omega) = \frac{1}{I_1(\omega)/I_2(\omega) + 1}$$

A slight complication is added due to the inhomogeneous broadening of the ground state separations, but will be necessary to explain experiments in chapter 5. For simplicity, we'll consider that state 2 is broadened relative to state 1, and introduce  $\delta$  to label an ion's ground state detuning.

$R_1(\omega, \delta)$  for a given ground state detuning  $\delta$  is the same as above, but the equilibrium

condition is different. We must now equate  $R_1(\omega, \delta)$  with  $R_2$  evaluated at  $\omega + \delta$ . Explicitly,

$$P_1(\omega, \delta) = 1/(1 + I_1(\omega, \delta)/I_2(\omega + \delta, \delta))$$

The total population in state 1 is found by summing over the inhomogeneous profile  $G(\delta)$ :

$$P_1 = \int_{-\infty}^{\infty} G(\delta)P_2(\omega, \delta)d\delta \quad (3.1)$$

If  $G(\delta)$  is a Gaussian (a good approximation for strain broadening), this becomes:

$$P_{in} = \int_{-\infty}^{\infty} \exp(-\delta^2/(2\sigma_s)) \frac{I_2(\omega + \delta, \delta)}{I_2(\omega + \delta, \delta) + I_1(\omega, \delta)} d\delta \quad (3.2)$$

### 3.3.6 Directly Addressing Spin Transitions (Using RF)

A final tool of some use in optical hole-burning is directly manipulating the ground states, using radio-frequency (or possibly microwaves for electronic Zeeman levels) pulses. This is the domain of solid-state NMR (or Electron Spin Resonance (ESR)).

Common uses in hole-burning are as a ‘reset’ where RF power is applied on all transitions to scramble their populations. More crucially, RF may be used to avoid slow hole-burning as may be caused by unfavourable branching ratios (i.e. the ions mainly relax back into their original state). By applying RF at the frequency of an excited state spin transition, a better branching ratio might be achieved.

RF is interesting for purposes where it might be used because it can be highly accurate. There are no propagation effects (the wavelength is much larger than typical sample sizes) and sources are available with high precision and stability.

### 3.3.7 Limitations of Hole-burning

Using hole-burning, it is possible to make arbitrary spectral structures as long as they are within the limits set by the material. The limits are on the maximum optical depth, narrowest spectral edge, and widest bandwidth. The material properties which ultimately limit these are the combination of oscillator strength and spectral density, the homogeneous line-width, and the spacings between the shelving levels.

The optical depth of a material is a product of the transition’s oscillator strength and number density of ions, and is divided by the inhomogeneous broadening. Assuming a situation as in Fig. 3.1, the relevant oscillator strength is that for the optical transition as a whole, as an equal number of ions absorb on transitions from each ground state. Using hole-burning we can theoretically increase the peak optical depth by the number of ground states: that is  $N$  ground states, an  $N$ -fold increase is possible. This is because any transition which may absorb at a given frequency has its ground state only  $1/N$  populated at equilibrium. By careful hole-burning, every ion which may absorb can be made to do so.

The width of the narrowest spectral absorption feature is limited by the homogeneous line-width to  $2\Gamma_h$ . The factor of 2 is because we are using light to create a population density, the spectrum of which is itself limited by  $\Gamma_h$ .

The bandwidth of any created structure is limited by the spacings of the shelving levels. This is particularly true for hole-burning with quadrupole levels where there may be several nearby levels. Which splittings are important depends on how many levels there are and relative size of the excited and ground splittings. This is an important limit for hole-burning memories and is worth some consideration. We will assume that the absorption on no transition can be neglected. This will be the case if we wish to make very optically thick structures.

If the excited state splittings are small, they will limit the free bandwidth around a feature, due to additional transitions of the ions in the feature. If the ground state splittings are small, the free spectral region is limited by their sum as this is the total frequency that an ion in the trench can be moved. In general it's a combination of both.

The situation shown in Fig. 3.1 serves as a relevant example (it corresponds to  $\text{Pr}^{3+}$  in a low symmetry site). The levels are split in three and the excited state splittings are smaller than those in the ground state. In this case the widest trench that can be made is the difference between the total splittings in the ground and the excited states,  $\Delta_{g,1} + \Delta_{g,2} - \Delta_{e,1} - \Delta_{e,2}$ . The ground states let us address any transition in our trench and move it, while the excited states mean we have to move them further.

The largest transmissive bandwidth around the feature is a little more complex. It is the same as the trench without the feature in it, *if* we don't require the feature at the centre of the trench. Let's say our feature consists of a single transition to the highest excited state. It will also be resonant at detuning lower by  $\Delta_{e,2}$ . We can avoid additional absorption anywhere else in the trench if we put our feature  $\Delta_{e,2}$  from the lower frequency wall of the trench.

We should also note the possibility to make repetitive wide-band structures, although they do not meet our requirements. To create the structure, the level splittings in the ground and excited states are carefully tuned to be equal. This method was used in recent demonstrations of a wide bandwidth AFC in  $\text{Tm}^{3+}[2]$ , using a magnetic field to tune the splittings. For memory purposes this structure is limited as it can't be used to achieve high efficiencies, doesn't leave a free ground state for long-term shelving, and can't be directly used to give a controllable delay.

A final limitation of hole-burning is the contrast that may be achieved between peak absorption and peak transmission. This is determined by whether we can make the structure faster than it collapses. To simplify, we assume perfect experimental hole-burning techniques which avoid propagation effects and laser line-width, and further neglect branching ratios and other factors slowing the decay (which can be avoided to some extent). In this case, the quickest we can move ions incoherently is  $T_1^{opt}$ , while the rate they return of their own accord is  $T_1^s$ . Theoretically we might use coherent techniques to transfer population

faster than this, but causing faster transitions means we have to affect larger bandwidths. The trade-off for rapid feature creation then is broader features.

### 3.4 Specific Materials for Quantum Memory

We now look at the choice of specific REI doped material for implementing a hole-burning memory. Various hole-burning materials have been studied for decades. Until recently however, the main goals have been applications such as lasers, Time-domain Optical Memory (TDOM), and fundamental spectroscopy. These applications, and the availability of applicable lasers for the research, is reflected in the type of materials most studied. Relatively recent material reviews have also been made[104, 109, 125], with a focus on classical TDOM. The requirements of quantum memory are subtly different. Reviews focussing on quantum applications (e.g. [25, 126]) have also been attempted, but the author feels they are incomplete. In particular, issues stemming from the fundamental limitations of hole-burning are largely un-addressed.

The aim of this section then is to very briefly outline the specific requirements of a hole-burning material for quantum memory, to outline the kinds of materials that meet these requirements, and to indicate the best materials presently available. This will lead us to choose a material for experiments in chapter 5.

#### 3.4.1 Target Material Parameters

As described in chapter 2, the combination of a high efficiency with a large TBP will be required for most applications. A memory based on hole-burning will be governed by the details of the spectral structure that can be created, and these are in turn limited by the material parameters(section 3.3.7). For applications such as quantum repeaters, an added requirement is the potential for long-term shelving. For this purpose an additional high-coherence shelving level is also needed.

In the first instance, the achievable efficiency and TBP will be limited by the quality of hole-burning. If we also assume that better than 50% efficiency is required, a much stronger requirement is placed on the spacing of the hole-burning level. These two factors are:

- $C = T_1^s/T_1^{opt}$ , the ratio of the speed at which the structure collapses to the speed with which we can create it, giving the contrast ratio between the height of the most absorbing feature to the residual absorption level. It gives the maximum broadening or finesse factor which a CRIB or AFC memory can use before the desired absorption is comparable to the background, and is thus proportional to the maximum TBP for a given efficiency.
- $(\mathcal{TB})_{max} = \min(\Delta_s)/(2\pi\Gamma_2^{opt})$  is roughly the ratio of the widest spectral trench to the narrowest absorption feature that can be created within it. A higher TBP can't

be realised without being limited to  $\approx 50\%$  efficiency from re-absorption.

For an application such as a repeater which will likely require a storage time of  $> 1s$ , the two additional requirements are

- Three ground states (1 for the original transition, 1 for hole-burning, 1 for shelving coherence)
- $T_2^s > 1s$ , the shelving state must have a sufficiently long coherence.

Finally we note two further practical requirements:

- The optical density determines how long the medium must be to achieve the required absorption, unless a multi-pass cells or cavity is used.
- For CRIB based memories, an appropriate Stark (or Zeeman, or other) shift may be a limiting factor. Further complications arise due to symmetry, as different orientations will shift differently in the presence of external fields.

These are aimed at quantum memory. The base requirements of classical TDOM are considerably less. This is because of the reduced requirement on noise and efficiency. For classical purposes the entire inhomogeneous line may be used, and the usual figure of merit is  $\Gamma_{ih}^{opt}/\Gamma_h^{opt}$ . We will discuss alternative methods to use the entire inhomogeneous line for quantum memory in chapter 6.

### 3.4.2 High-Coherence Transitions

Considering the decades of research into hole-burning materials, a thorough review is far beyond our present scope. We note instead that to date the transitions that have mainly been studied for QI applications (although certainly not the only ones) are the  ${}^3H_4 \rightarrow {}^1D_2$  transition of  $Pr^{3+}$ , the  ${}^7F_0 \rightarrow {}^5D_0$  transition of  $Eu^{3+}$ , the  ${}^3H_6 \rightarrow {}^3H_4$  transition of  $Tm^{3+}$ , and the  ${}^4I_{15/2} \rightarrow {}^4I_{13/2}$  transition of  $Er^{3+}$ . The best coherence properties have been observed when doped into the relatively recent host material,  $Y_2SiO_5$  (YSO), with the exception of  $Tm^{3+}$  which has better reported properties in  $Y_3Al_5O_{12}$  (YAG). We summarise the relevant properties briefly (some further details of YSO are given in section 3.4.4).

#### ${}^4I_{15/2} \rightarrow {}^4I_{13/2}$ transition of $Er^{3+}$

Starting at the right of the lanthanide group, the  $1.5\mu m$   ${}^4I_{15/2} \rightarrow {}^4I_{13/2}$  transition of  $Er^{3+}$  is of great technological importance in telecommunications. It is a Kramers ion with a corresponding strong magnetic character, making it unfavourable in the first instance for achieving high coherence. In YSO (as with other materials) in standard conditions of  $\approx 4K$  and low fields, it decoheres rapidly. It has on the other hand an unusually long lifetime of several milliseconds.



Its coherence properties are greatly improved in the presence of a strong magnetic field ( $\approx 7\text{T}$ ) and lower temperatures[127]. This has the twin effect of helping to reduce dephasing by lattice spin-flips and increasing the electronic Zeeman splitting to greater than  $kT$ . The spin levels then relax thermally to a single ground state[106] so that cross-relaxations between  $\text{Er}^{3+}$  ions are avoided. In this situation,  $\text{Er}^{3+}:\text{Y}_2\text{SiO}_5$  (ErYSO) has been used to demonstrate the narrowest homogeneous line-width yet reported in a solid at  $73\text{Hz}$ [127] ( $T_2 = 4.4\text{ms}$ ). This is still less than the lifetime limit, The lifetime is also increased by the field and temperature to as high as  $T_1 = 11\text{ms}$ .

Hole-burning in ErYSO has been demonstrated using electronic Zeeman levels. At the high fields used to obtain good optical coherence properties, a lifetime of  $40\text{ms}$  was measured for the spin states[106]. This was attributed to the direct phonon process which increases with field intensity. Investigations intent on maximising this at moderate fields found a dominant lifetime of  $130\text{ms}$  [128]. A small level of longer-lived population ( $60\text{s}$ ) was also noted, thought to be due to super-hyperfine structure.

Using Zeeman levels for hole-burning, the contrast ratio that can be expected is only  $C \sim 10$ . This is the major limiting factor. Due to the large Zeeman splitting at large fields  $\sim\text{GHz}$ , the available TBP is not a large problem with  $(\mathcal{TB})_{max} \sim 10^7$ .

### ${}^3H_6 \rightarrow {}^3H_4$ Transition of $\text{Tm}^{3+}$

Moving lower on the periodic table, the  ${}^3H_6 \rightarrow {}^3H_4$  transition of  $\text{Tm}^{3+}$  usually occurs around  $800\text{nm}$ . This corresponds to an earlier telecommunications window and remains convenient for compatibility with other systems.  $\text{Tm}^{3+}$  is a little unusual as its optical coherence properties tend to be better in garnets than YSO, most likely due to a phonon dephasing from a low-lying crystal-field level. The best coherence properties have been reported in YAG[125]. It has lower magnetic sensitivity than  $\text{Er}^{3+}$  due to being non-Kramers. However, while relatively long lifetimes  $T_1 = 2\text{ms}$  have been reported, the highest coherence time is much lower, with  $T_2 \approx 500\mu\text{s}$  in a magnetic field[125]. The dephasing is attributed partly to phonon broadening from the nearby crystal-field level, and to a relatively large nuclear magnetic moment due to an enhancement by interaction with its electrons.

For hole-burning,  $\text{Tm}^{3+}$  is an interesting case. It has a nucleus with  $I = 1/2$ , giving it no quadrupole levels and no hyperfine structure in zero field. This can greatly simplify two-level experiments, but limits hole-burning experiments or those requiring long-lived coherence. Hole-burning has been demonstrated both using a metastable  $LS$  level[129] and using nuclear Zeeman levels. The  $LS$  level has the label ( ${}^3F_4$ ), which sits in one decay path ( ${}^3H_6 \rightarrow {}^3H_5 \rightarrow {}^3F_4 \rightarrow {}^3H_4$ ) from the normal excited state to the ground. It is only useful for short-term hole-burning due to its lifetime of  $\approx 10\text{ms}$ .

Hole-burning using the nuclear Zeeman splitting is a much better prospect for long-lived hole-burning. Either a specific magnetic field or a using RF to mix the excited states will be required to get a good branching ratio. Hole-burning was demonstrated with a

lifetime of  $\approx 30$ s by Ohlsson et al.[108].

Using Zeeman states for hole-burning, relatively large TBPs are theoretically possible, depending on the field applied. Splittings of 100MHz or greater are observed, so we may estimate  $(\mathcal{TB})_{max} \sim 10^5$ . The hole-burning contrast ratio is the major limiting factor at  $C \sim 1 \times 10^3$  (or  $\approx 10$  depending if we consider the Zeeman state or the metastable state). Using Zeeman levels for hole-burning however, we cannot apply long-term shelving techniques as there are no free ground states.

#### ${}^7F_0 \rightarrow {}^5D_0$ Transition of $\text{Eu}^{3+}$

The  ${}^7F_0 \rightarrow {}^5D_0$  transition of  $\text{Eu}^{3+}$  occurs near 580nm. It displays the most extreme properties of the lanthanides, at least partly because both the excited and ground states have no magnetic moment, that is they have  $J = 0$ . This makes it the only REI with a true singlet ground-state, independent of the host symmetry. The transition strength is usually quite weak but is very host dependant.

In  $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$  (EuYSO), the optical transition has a lifetime of  $T_1 = 2$ ms in zero field and a coherence time of 1ms. Lifetime and coherence times are improved in a small magnetic field ( $\sim 100$ G) to reduce Y-Y spin flips. In this case line-widths as low as  $1/(\pi T_2) = 105$ Hz ( $T_2 = 3$ ms) have been estimated, very close to the lifetime limit[130]. The oscillator strength is  $\sim 1 \times 10^{-8}$ .

EuYSO also has very good demonstrated hole-burning properties again due to the low magnetic sensitivity of the ground state. At zero field it has nuclear quadrupole levels due to its  $I = 5/2$  nucleus. A peculiarity of Europium ion is that the nuclear quadrupole splittings are enhanced by interaction with electrons ( $\sim 50$ MHz in YSO) while the Zeeman shift is reduced[131, 111]. The quadrupole levels have been estimated to last upwards of 20 days[110], with a minimum of spectral diffusion[132]. This is likely related to the Relatively long spin coherence times of 15ms[133] were measured in zero field, increasing to 36ms with 100G applied. The decay with the field applied was non-exponential and was attributed to Y-Y spin flips.

Studies of the spin Hamiltonian indicated[134] that ZFOZ techniques will be effective to reduce Y-Y dephasing and increase coherence times to beyond 100s. One difficulty is the requirement of relatively large fields ( $\approx 2$ T), and cryostats that don't leak<sup>2</sup>.

In YSO it should achieve a very large contrast ratio using hole-burning of  $C \sim 1 \times 10^8$ . The maximum efficient TBP is roughly  $(\mathcal{TB})_{max} \sim 1 \times 10^5$ .

#### ${}^3H_4 \rightarrow {}^1D_2$ Transition of $\text{Pr}^{3+}$

The  ${}^3H_4 \rightarrow {}^1D_2$  transition of  $\text{Pr}^{3+}$  occurs near 600nm. It is well studied due to the early availability of Rhodamine-6G dye lasers for accessing its wavelength. It has good demonstrated optical coherence and hole-burning properties. In comparison with  $\text{Eu}^{3+}$

<sup>2</sup>The author may have been involved in such failed experiments

its coherence and lifetime are worse but its oscillator strength is larger. It is also more studied.

In YSO,  $\text{Pr}^{3+}$  has fair optical lifetime of  $160\mu\text{s}$ , and a line-width of  $1.7\text{kHz}$  ( $T_2 = 190\mu\text{s}$ )[135] in zero field. Its oscillator strength was reported by Equal et al. as  $3 \times 10^{-7}$

For hole-burning  $\text{Pr}^{3+}$  has a  $I = 5/2$  nucleus similar to  $\text{Eu}^{3+}$ . The corresponding nuclear quadrupole levels have been used successfully for hole-burning on the scale of several minutes. The quadrupole splittings however are relatively small,  $\sim 10\text{MHz}$ .

Using the quadrupole levels, the best contrast ratio is  $C \sim 10^6$  while the maximum efficient TBP is  $(\mathcal{TB})_{max} \sim 10^3$ .

To date, the only demonstration long-term shelving at the appropriate level was also performed in PrYSO. This used a combination of DDC and ZFOZ of 30s[122].

For the discussed properties, the material with the best long-term properties is EuYSO. Its primary issue is a practical one, that of relatively low oscillator strength. An easier target for a demonstration is PrYSO.

### 3.4.3 Future Materials

We may now make some general comment on material requirements and limitations. The aim is for a high coherence optical transition, good hole-burning, and long-lived shelving.

The main factors limiting optical coherence are phonon interactions with the host, interactions between dopants, and magnetic dephasing by the host. We start by ruling out materials with inherent dephasing such as glasses or those with mobile charges.

The transitions of a specific ion that couple least to phonons are those from the lowest level of a  $J$ -multiplet to the ground state. In a good material, most of the oscillator strength of such a transition should be in the zero-phonon line. This implies an asymmetric dopant site, so that the lowest-lowest transition is electric-dipole allowed, and that the  $J$  levels are well separated.

A material should be relatively lightly doped to avoid interactions between dopants. A highly ordered host material is desirable for this reason, to avoid inhomogeneous broadening. While desirable in classical TDOM applications, inhomogeneous broadening reduces the available optical depth. The concentration may be increased to compensate, but at the expense of greater dephasing between dopants. In some cases, material disorder may cause dephasing directly due to charge traps or other defects mediating dephasing interactions.

In keeping with low disorder, it is better for the dopant to be at a substitutional site rather than occupying an interstitial or other defect site. In the first place, this reduces the induced crystal strain. In the second, a single type of defect site rarely exists in isolation. Thus it is likely that the dopants will also occupy other sites with different optical properties. These extra dopants will cause broadening and spin dephasing without contributing to the absorption. Similarly the dopant site should be a good match both for the size of the ion, and for its  $3+$  valence state.

The final dephasing mechanisms may be magnetic interactions by spins in the host.

Mechanisms exist for reducing these (section 3.2.10), but may be partially avoided in the first place by choosing a host with a low concentration of spins.

The limiting factor for optical coherence Kramers ions such as  $\text{Er}^{3+}$ ,  $\text{Nd}^{3+}$ , or  $\text{Dy}^{3+}$  will likely be interaction between dopants, unless they may be spin-polarised. For non-Kramers ions like  $\text{Tm}^{3+}$ ,  $\text{Pr}^{3+}$  and  $\text{Eu}^{3+}$ , the relevant transitions may be limited by the radiative lifetime even in low fields. In a good host-ion combination, the remaining non-lifetime dephasing may be due to magnetic interactions, which may in principle be reduced as in section 3.2.10.

Of the hole-burning mechanisms discussed, those which offer the best hole-burning are the nuclear spin levels, in particular those due to the nuclear quadrupole interaction. Memories based on hole-burning using these levels may instead be limited in bandwidth by their splitting. Hole-burning using electronic Zeeman or other optical transitions do not have the same bandwidth limitations, but the limitation of contrast ratio seems greater than the limitation of bandwidth for quadrupole levels.

For the purposes of long-term storage, at present the only method with the demonstrated performance is using ZFOZ techniques on a quadrupole transition as by Fraval et al.[122]. It may be plausible to achieve similar times with nuclear Zeeman levels such as  $\text{Tm}^{3+}$ . Long term coherence on ( $I = 1/2$ ) nuclear Zeeman levels has been achieved in a very dilute system of a phosphor-in-silicon[91] using DDC techniques. The limiting factor for  $\text{Tm}^{3+}$  will be coupling between dopants in the (relatively) concentrated rare-earth material. For electronic Zeeman levels, this difficulty will be many times greater. Further, the procedure to obtain good optical coherence in high fields is incompatible with good spin-shelving.

Using present techniques then, a Kramers ion material cannot meet our requirements for a hole-burning memory. Of the non-Kramers ions discussed,  $\text{Eu}^{3+}$  is the major standout followed by  $\text{Pr}^{3+}$ . The other non-Kramers rare-earths appear to have other issues.  $\text{Pm}^{3+}$  is radioactive with a very short half-life.  $\text{Tb}^{3+}$  and  $\text{Ho}^{3+}$  both have nuclear quadrupole moments, but they appear to be less studied for information storage purposes.  $\text{Tb}^{3+}$  tends to suffer phonon dephasing from a near degenerate crystal field splitting in the ground state, while  $\text{Ho}^{3+}$  suffers magnetic dephasing due to an exceptionally large enhanced nuclear magnetic moment[100].

#### 3.4.4 Further details of PrYSO

Given presently available materials, we choose PrYSO for experiments in chapter 5. This is due to its conveniently large oscillator strength, long hole-burning lifetime, and the degree to which its properties are known. As noted above, it is more limited for long-term applications than materials like EuYSO due to smaller hyperfine splittings, but is appropriate for a demonstration. It also has a solid record for quantum memory techniques with EIT for over 2s[1] and a proof-of-principle demonstration of a gradient echo using the Stark effect[136]. We briefly review some additional properties.

There are two distinct Yttrium sites in YSO, usually labelled imaginatively as Site 1 and Site 2, each with a distinct  $C_1$  symmetry.  $\text{Pr}^{3+}$  ions may substitute at either, but with a significant preference at “Site 1”. The other, “Site 2” has a slightly longer lifetime ( $T_1 = 220\mu\text{s}$  vs  $T_1 = 160\mu\text{s}$ )[130] but also smaller hyperfine splittings of only a few MHz. The 2 sites have roughly the same oscillator strength, with “Site 1” being much more absorbing due to the greater occupation. The longer lifetime may also be related to the lower occupation of “Site 2”. We are concerned with “Site 1” because of the much greater optical depth and larger hyperfine splittings. At zero field “Site 1” has been estimated to have a homogeneous line-width of 2.4kHz by Equall et al.[130], although Graf et al. [135] measured it as 1.7kHz. This indicates the dependency on crystal quality, as the dephasing is thought to be largely due to Y-Y spin flips. The lifetime limit is 1.0kHz.

Pr has only one naturally occurring isotope, with a nuclear spin of 5/2. This leads to 3 ground nuclear quadrupole levels in zero field. Transitions between all levels occur to some extent, but mixing is strongest between 1/2 and 3/2 levels. The 5/2 levels form a somewhat isolated system, but decays to the other ground states occur  $\sim 1\%$  of the time[137]. Decays into and out of this system are thus still useful for hole-burning.

The lifetime of the ground states, and thus of spectral holes, are measured in minutes (we make a measurement in chapter 5). The Nuclear electric-quadrupole and Zeeman tensors which determine the hyperfine splittings have been completely determined [138, 139]. This has enabled the implementation of high precision decoherence control techniques[122, 140].

The optical depth of PrYSO crystal varies somewhat with doping concentration, but is generally on the order of 100dB/cm. The inhomogeneous line-width is around 10GHz, but varies very strongly with doping concentration(Fig. 3.5).

The YSO crystal has  $C_{2h}$  symmetry. This means there are 4 orientations where  $\text{Pr}^{3+}$  can substitute, and thus 4 different dipole moments, arranged in a kind of X pattern. This means that for performing a gradient echo using Stark shifts, it will be necessary to select dipole orientations using hole-burning. The Stark shift is 112kHz  $\text{V}^{-1}\text{cm}$  [141].

## 3.5 Alternatives to Metastable Hole-burning

### 3.5.1 Persistent Hole-burning

In addition to hole-burning using metastable levels, several mechanisms have been found to create semi-permanent holes. The general problem is to convert the short-lived optical excitation of an ion into a persistent change of state or its surrounding environment. Divalent  $\text{Eu}^{2+}$  and  $\text{Sm}^{2+}$  when doped into some materials may be photo-ionized in so-called photon-gated hole-burning[142, 143]. After excitation, an additional beam may be used to excite the electrons excited into the conduction band to be stored elsewhere. In hydrogenated  $\text{CaF}_2$ , hydrogen ions may occupy different configurations of interstitial sites near REI dopants to provide charge compensation. Optically exciting the REIs can trigger

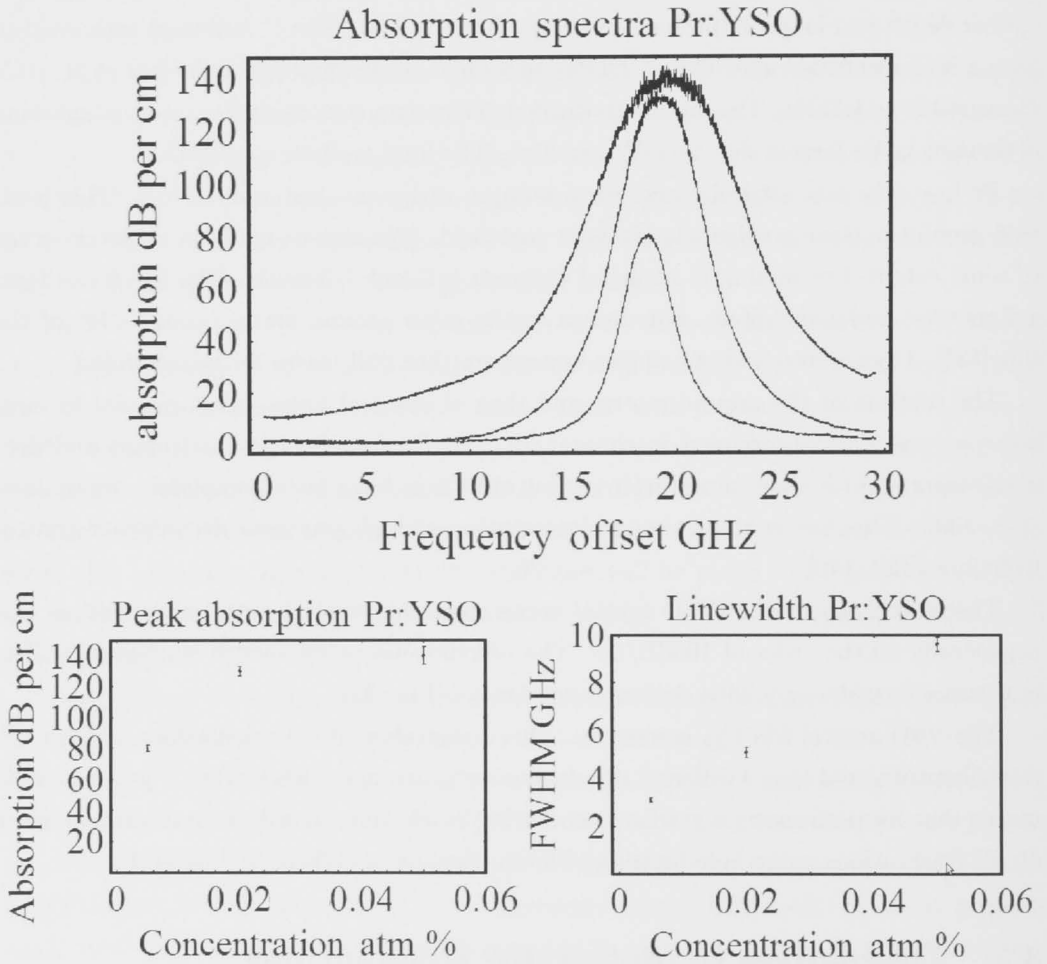


Figure 3.5: The inhomogeneously broadened lines of 3 different concentrations of Pr:YSO. Measurements are due to Annabelle Alexander, based on transmission along the  $C_2$  axis of 1mm samples from Scientific Materials.

the configuration of the hydrogen ions to change, altering the field seen by the REI once it relaxes again.

While such mechanisms are on one hand attractive due to their high stability, they generally cause large changes to the medium. This means that more than just the targeted ions are effected. The spectral features that can be created using them tend to be quite broad, and broaden rapidly as the depth is increased. They are thus unlikely to be appropriate for quantum memory.

### **3.5.2 Laser Induced transition**

A different approach to achieving narrow spectra, with some success, is using auxiliary light fields. Both holes and anti-holes can be created using an additional beam and a 3 (or more) level system, as described in chapter 2. The requirements for creating a Raman anti-hole are almost identical. In the limit of large detuning, the width of the created anti-hole is dependant on the laser and the width of the ground state transition, rather than the width of the optical transition. The issue is that in past demonstrations it's been necessary to first perform hole-burning to create a free ground state. Thus the limitations on bandwidth are the same.

In specific cases it should be possible to remove the bandwidth limitations. We'll discuss this further in chapter 6.

### **3.5.3 Strain-free Crystals**

Perhaps the most obvious way to defeat inhomogeneous broadening is to avoid it in the first place. If the broadening is reduced enough, a crystal in its equilibrium state would have permanent, narrow features corresponding to its hyperfine structure. While relatively narrow inhomogeneous line-widths have been observed in the past[144, 145], this has only been at very low dopant-concentration in high symmetry sites, giving very weak absorption. Aside from this limitation, the emphasis has largely been on large inhomogeneous broadening for classical information applications.

A slightly different direction for achieving low broadening is using stoichiometric materials, which we will discuss in 6.2.4. This has the immediate obvious advantages of much greater optical depth. We'll also see that if the broadening can be reduced below the level of hyperfine structure, it may also remove bandwidth limitations. To see this we must first discuss properties of the gradient-echo technique in the next chapter.

### **3.5.4 Variations of a 2-pulse Photon Echo**

A final possibility is to avoid the need for sharp structure. This is the case for a standard 2PE. The possibility of using a variant of the 2PE for quantum memory has been revived recently[66, 146], where an echo in the excited state is suppressed. These techniques do

---

not require hole-burning, and are thus not necessarily limited by hyperfine structure. We discuss them further in 6.2.2.





---

# Optical Gradient-Echo - Theory and Modelling

---

We now focus on a ‘gradient echo’ memory technique, a specific form of CRIB introduced in chapter 2. It is peculiar amongst the techniques discussed in chapter 2 in that it involves a spatial variation in the storage medium. This creates complications, but also offers advantages. One consequence is that it is not necessary to reverse the retrieval direction to avoid reabsorption, and to thus obtain high efficiencies. This significantly lowers requirements for a high efficiency demonstration, and is the reason we use it for experiments presented in the next chapter. Before this, it will be helpful to study the technique in some detail, away from the complications of the lab.

The base aims of this chapter are to explain how and why the technique works, and to give the detail necessary to interpret our experimental results. A broader aim is to describe further properties of the technique which will motivate future experiments. We’ll find that aside from being relatively simple to implement, variations of the technique may be used to powerfully manipulate light modes for use in repeaters and other applications. Further, its experimental requirements are fewer than previously thought, and in the long run may solve the problem of limited bandwidth due to hyperfine structure of memories using REIs.

We’ll broadly use a numerical model for our study. The main variation from previously studied models is that we include inhomogeneous broadening of the initial absorption feature.

After specifying our model, the flow of the chapter is as follows.

In section 4.2 we will consider the technique as originally described[67], without variation. This means we begin with a narrow, homogeneously broadened feature and control its evolution by inducing a detuning that changes linearly with depth. We’ll describe the space-dependent response of the medium and how it affects propagation. We’ll consider several ways of looking at the storage, in time, frequency, and k-space.

In section 4.3 we’ll look at variations on the initial feature shape, specifically those whose shape is mainly dependent on inhomogeneous broadening. We will find that as the optical depth is increased, one side of the initial spectrum becomes more significant. This

leads to a sharp spectral edge being required for an efficient memory, rather than a narrow line. It further introduces the possibility of hiding unwanted spectral structure.

In section 4.4 we will consider generalisations of the detuning field. This field allows manipulating the resonant frequency and phase of the absorbers very precisely, making it a very direct and powerful tool for manipulating the light which is mapped onto them. We also confirm that the technique is quite robust to experimental imperfections.

Section 4.5 gives a very brief discussion on how to optimise the memory efficiency. We do not solve the problem, but illustrate the difficulty of simultaneously optimising in both time and frequency space. This is further complicated by propagation effects which mix the two. We consider the dependence on feature shape in a specific case, and describe a method to find a pulse shape that fits a given memory.

We finish with a brief summary of experimental requirements for a high efficiency demonstration.

## 4.1 Gradient-Echo Model

We model the gradient echo we using an approximate form of the Maxwell-Bloch equations (section 1.3.6). Previous work has also used these[136, 58, 147, 148, 149], however to the best of the author's knowledge none have considered intrinsic inhomogeneous broadening of the initial ensemble, only considering one broadened by homogeneous decay.

The general problem we consider is a medium extending from  $-z_0$  to  $z_0$ , populated by a uniform density of absorbers. The spectral density of the absorbers is described by  $N(\Delta)$ . A space dependent external field is applied which shifts the resonant frequency of every absorber by  $\Delta_{ext}(z)$ .

The polarisation of each absorber is then driven by the incident field,  $E$ . We use equations for the slowly-varying envelope of each quantity (sections 1.3.4 and 1.3.6). A single atom has polarisation  $\alpha$  described by

$$\dot{\alpha}(t, z, \Delta) = -\left[\frac{\gamma}{2} + i(\Delta + \Delta_{ext})\right]\alpha(t, z, \Delta) - iw g E(t, z) \quad (4.1)$$

The excitation of the single atom,  $w$ , evolves as

$$\dot{w} = \Im(\alpha E) - \gamma(w + 1) \quad (4.2)$$

The polarisation density of the medium,  $P$  is sum over atoms resonant at all detunings:

$$P = i \int_{-\infty}^{\infty} N(\Delta) \alpha(t, z, \Delta) d\Delta \quad (4.3)$$

The propagation of the electric field is described (in the co-moving frame) as

$$\frac{\partial}{\partial z} E(t, z) = iP \quad (4.4)$$

where we are using units such that  $\frac{k_0}{2\epsilon} = 1$ .  $iP$  is the source term describing the (re)radiation field due to  $P$ .

The standard ‘gradient’ echo uses a detuning gradient such that  $\Delta_{ext} = Gz$ , with  $G$  having units of frequency per unit length. Further it is normally assumed that the medium in the absence of the field is purely homogeneously broadened by the decay rate  $\gamma$ , giving a homogeneous (Lorentzian) line-width  $\Gamma = \gamma/(2\pi)$ . That is,  $N(\Delta) = \delta(0)$ . This is the situation we consider first, before generalising in sections 4.3 and 4.4.

We only consider the weak pulse regime such that eq. 4.2 is  $w \approx -1$ . Nevertheless, deviations above  $w = -1$  are a useful indicator of the energy stored in the atoms at any instant.

These equations are simple to solve numerically due to their slowly varying, cross-propagating form. Simulations in this thesis used a C++ program which was generated mostly automatically using an open-source code-generator, XMDS[150]. The generated program used a 4th order Runge-Kutta method, solving for  $E$ ,  $\alpha$ , and  $w$  for a 3D grid of  $z$ ,  $t$ , and  $\Delta$ .

#### 4.1.1 Units and Simulation Parameters

Finally it’s worth commenting on units. In the above equations, the length and time scales of the wavelength and wave period (and also the speed of light) have been factored out. The only relevant length scale is that of the sample ( $2z_0$ ), whereas we have no time reference. As units of time and frequency we will tend to choose  $\mu\text{s}$  and MHz in order to be similar to our experiments next chapter, but the choice is arbitrary.

The major parameters governing a given simulated memory are those discussed in chapter 2. These are the memory bandwidth (given by  $2G_0$ ), the decay rate of a single absorber ( $\gamma$ ), the inhomogeneous spectrum of the initial feature ( $N(\Delta)$ ) and the optical depth of the medium before and after the broadening is applied.

## 4.2 Unmodified Optical Gradient-Echoes

Before we consider variations, we look at the original technique. Specifically, we start with a narrow absorption feature defined by a transition with dephasing rate  $\gamma/2$ . A linearly increasing (with depth) detuning field  $\Delta_{ext} = Gz$  is applied to broaden it. The low optical-depth may be understood as any other echo: spectral components are absorbed resonantly, the relative phase of the ensemble is then evolved backwards until they re-enact the absorption process. In this case the spatial dependence is irrelevant, but it becomes important as the optical depth is increased. In the first instance, it is responsible for the echo not being reabsorbed as it leaves the medium.

This situation was studied analytically by Longdell et al.[58], and by Moiseev and Arslanov[151]. Longdell et al. used a symmetry argument to estimate the output of the memory in the regime of high optical depth, no decay ( $\gamma = 0$ ), and much larger memory

bandwidth than pulse bandwidth. An analogy with beam-splitters was used to obtain the efficiency when the optical depth was lower. Moiseev and Arslanov also included homogeneous decay in their analysis, and obtained a more precise (although significantly more complex) expression for the output field in terms of the input. We will take the slightly less general, but much simpler approach.

Longdell et al. assumed that the field gradient was flipped instantaneously at  $t = 0$ . They found that when the optical depth of the medium is large, the output field is given by

$$E_{out}(t) = -E_{in}(-t)|t|^{2i\Phi} \frac{\Gamma(i\Phi)}{\Gamma(-i\Phi)} \quad (4.5)$$

where  $\Gamma$  is the gamma function, and  $\beta = 2\pi\Phi$  is the peak optical depth of the broadened medium (the original unbroadened feature is infinite in height, as we're assuming it's infinitely narrow). The factor  $|t|^{2i\Phi} \frac{\Gamma(i\Phi)}{\Gamma(-i\Phi)}$  has a magnitude of 1, so that the output is a phase-modulated (as well as time-reversed) copy of the input. The phase of this modulation is

$$\delta\theta = 2\Phi \ln(t). \quad (4.6)$$

The modulation increases with the broadened optical depth, but decreases with longer storage times. They also found that as the optical depth is reduced, the efficiency is similar as for other memory techniques, given by  $(1 - e^{-2\Phi\pi})^2$ .

Building on this model, we may estimate the effects of homogeneous decay quite simply. Light that entered the medium at  $-t$  will be recalled approximately at a time  $t$ , giving a decay factor of  $e^{-\gamma t}$ . The exception is if  $\Phi$  is very large and  $t$  very small, so that the phase factor above will be associated with a delay. Our final approximate expression is

$$E_{out}(t) = -E_{in}(-t)|t|^{2i\Phi} \frac{\Gamma(i\Phi)}{\Gamma(-i\Phi)} (1 - e^{-2\Phi\pi})e^{-\gamma t} \quad (4.7)$$

which is valid for moderate  $\Phi$ , and  $Gz \gg \Delta\omega_{\text{pulse}} \gg \gamma$  ( $\Delta\omega_{\text{pulse}}$  is the pulse bandwidth).

In this section we will first try to give some insight into the above by looking at the space-dependent response of the medium (section 4.2.1). Following this we'll look at simulated examples of echoes. The normal situation of a time-domain memory will demonstrate the time-dependent properties, while using it as a frequency domain memory will illustrate frequency dependent properties. Finally a visualisation in spatial frequency ('k-space') allows the identification of a normal mode and gives insight into the medium's time evolution.

### 4.2.1 Space-dependent Response in a Gradient Echo

To account for the peculiar properties of a gradient echo, we must study its space dependence. To visualise how light propagates during the process, we simply add a space dimension to the standard absorption/dispersion plots (Fig. 4.1).

The response function for an initial Lorentzian profile of line-width  $\Gamma = 2\pi\gamma$ , broadened

by the detuning field  $\Delta = Gz$  is given by

$$\kappa = i\kappa_0 \frac{1}{(\omega \pm Gz) - i\gamma}$$

where we have rolled the normalising constants into  $\kappa_0$  (recall that  $\kappa$  simply describes an exponential change in the electric field envelope as it propagates, due only to the medium). The attenuation rate is  $-\Re(\kappa)$  (Fig. 4.1(a)), while  $\Im(\kappa)$  describes a change in phase.

If we consider a simple transmission measurement, the initial narrow feature appears to be smeared by the detuning gradient, although if we could apply the same measurement to a single slice of medium it would remain spectrally sharp.

We look at the propagation in two regimes. In the simplest case, the medium is optically thin, and/or the pulse is allowed to propagate for a long time. In this regime delays due to slow-light effects may be neglected. It is then only necessary to consider resonant absorption as represented by  $\Re(\kappa)$ . At the other extreme is short storage times and large optical depths, where slow-light effects are strong. In this case the light will never interact on resonance, and we only need consider dispersion ( $\Im(\kappa)$ ).

We start with the low OD, long storage-time regime. A wave-packet centred at any frequency will propagate until it is either through the medium, or it is absorbed resonantly (or its coherence/energy has been dissipated by some decay mechanism). Light centred at any frequency is absorbed by atoms resonant at that frequency.

Now when the detuning field is reversed, the standard echo picture as in 1.13 applies. The detunings of the atoms are reversed, and they rephase to the same configuration as during the initial absorption, and re-radiate. Each spectral component is free to propagate out of the medium without again encountering resonant absorption<sup>1</sup>. This is the reason why, unlike other echo techniques, we can have efficient retrieval in the forwards direction.

As the optical depth is increased and the propagation time is shortened, dispersive effects due  $\Im(\kappa)$  become significant. This because the light is still progressing in the medium at the time the detuning field is reversed.

We can gain some insight into this regime by considering the way a wave-packet is slowed it travels through the medium, as a function of its centre frequency and propagation depth. We look at the group-delay (section 1.3.8). We must be careful with interpretation however- we may only think about pure frequency components if we consider infinite times- but we are only considering until the field is reversed. Similarly a pure frequency has infinite length in our co-moving frame. Nevertheless, we may interpret the delay literally if we consider times short compared to the total group-delay, but long compared to the length of the wave packet. This may be achieved in a strongly broadened, high optical depth medium.

The group delay for the medium is calculated as the integrated delay from the beginning

---

<sup>1</sup>It may help to recall that we're in the co-moving frame where atoms with the same phase radiate constructively regardless of position.

of the medium at  $-z_0$  to  $z$ . We have:

$$D_g = c \frac{\partial}{\partial \omega} \int_{-z_0}^z \Im(\kappa) dz \quad (4.8)$$

where the integral is the total phase shift experienced by a wave packet centred at  $\omega$  after travelling to the point  $z$ .

Substituting for  $\kappa$ , we get

$$D_g = \frac{\kappa_0}{2G} \Re\left(-\frac{1}{\omega - Gz_0 + i\gamma} + \frac{1}{\omega + Gz + i\gamma}\right) \quad (4.9)$$

If we consider only frequencies well within the memory bandwidth, and also that the decay term is negligible, the first term above becomes small and we have just:

$$D_g = \kappa_0/2G \left(\frac{1}{\omega + Gz}\right) \quad (4.10)$$

We can consider a pulse whose bandwidth is well within the memory's, but is incident for a time  $T$  that is shorter than the maximum delay. Rearranging the above, the position of the pulse as a function of time is

$$z_g = \frac{\kappa_0}{G^2 T} + \frac{\omega}{G} \quad (4.11)$$

As it progresses in the medium, the group-delay it experiences increases (and the group-velocity decreases) as it is mapped onto atoms asymptotically closer to resonance.

As soon as the pulse is significantly delayed, it's net electric field is near zero, and it exists almost entirely as a propagating polarisation wave. When we manipulate the atoms by reversing the detuning field, we manipulate the pulse. The excitation is carried onto the other side of resonance, and the pulse is free to accelerate out of the medium. However, because the excitation hadn't yet reached resonance, the pulse's spectral components aren't flipped about zero. The shift they receive depends on their position in the medium- in this case, they get a shift of  $\delta\omega = G\delta z_g = \frac{\kappa_0}{GT}$ . As the pulse is retrieved at a later time again of  $t = T$ , we expected it to receive a phase shift due to this change in frequency of  $\int_0^t \delta\omega dt = \int_0^t \frac{\kappa_0}{GT} dt$ , so that  $\delta\theta = \frac{\kappa_0}{G} \ln(t)$ . This is the cause of the phase shift in eq. (4.7).

The situation where the broadening is comparable to the pulse bandwidth is more difficult. The group delay in both cases is given in Fig. 4.1. For the large bandwidth case, all spectral components approach resonance at the same rate. There will also be attenuation due to dephasing and transmission, but as the effect is the same for all frequencies, the output envelope is not distorted. In the low bandwidth case each frequency is different, in terms of the delay it receives and its level of attenuation. We'll look at this more practically using a numerical model in section 4.2.3.

The extreme of very high dispersion and a short storage time does not make for a good quantum memory, but is of interest. In this case the input pulse will be slowed and compressed into a thin slice near the input side of the medium (at the top of 4.1f). The

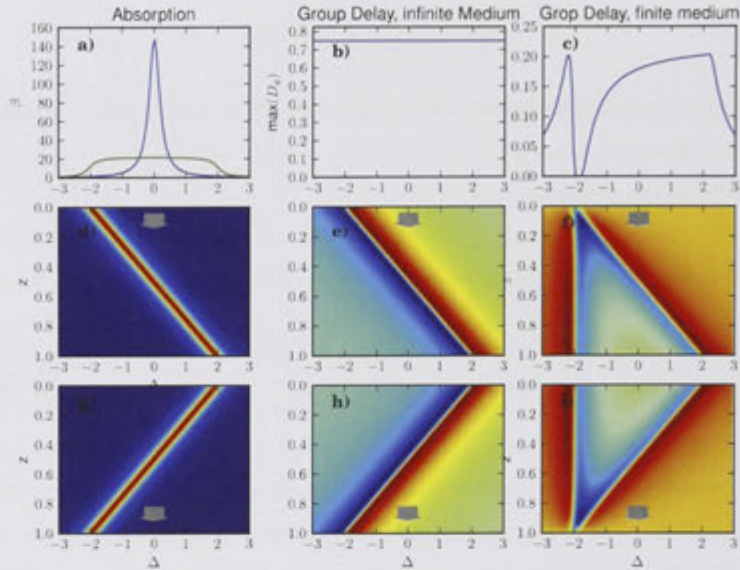


Figure 4.1: Absorption and group-delay due to an initial Lorentzian feature, broadened as in gradient echo. Top left shows total absorption of the sample, before and after it is broadened. The other two top panels show the maximum group delay experienced by a wave-packet in a medium with finite broadening, and with infinite broadening. The lower panels show the properties (absorption and group delay) as a function of detuning and position, before and after the field switch. The arrows indicate the propagation of a pulse.

propagation rate will be constant for all frequency components. When the field is switched, the components will be shifted in frequency by the same amount, and will otherwise continue to propagate at the same rate. When the pulse eventually exits the medium, it will be frequency and phase shifted, but otherwise just delayed, not temporally reversed. This gives us a way to think of the phase shift in this standard scheme as a competition between a simple delay + frequency shift, and a time-reversed storage and recall.

#### 4.2.2 A Time-domain Memory

We now illustrate the gradient echo with concrete (numerical) examples. In a practical application of a time-domain memory, relatively short pulses will be input, with bandwidths comparable to the level of broadening. Multiple inputs are separated in time.

Figures 4.2 and 4.3 show storage with high and low optical depths respectively. We can see in the low optical depth limit behaves similarly to other echo processes (Fig. 1.13), where the echo may be thought of as a repetition of the absorption process. As the optical depth is increased, the efficiency is increased accordingly. As it is increased further, the light is slowed by the large dispersion (Fig. 4.3a), and is still progressing in the medium at the time the field is switched. Associated with this is a phase shift and some distortion



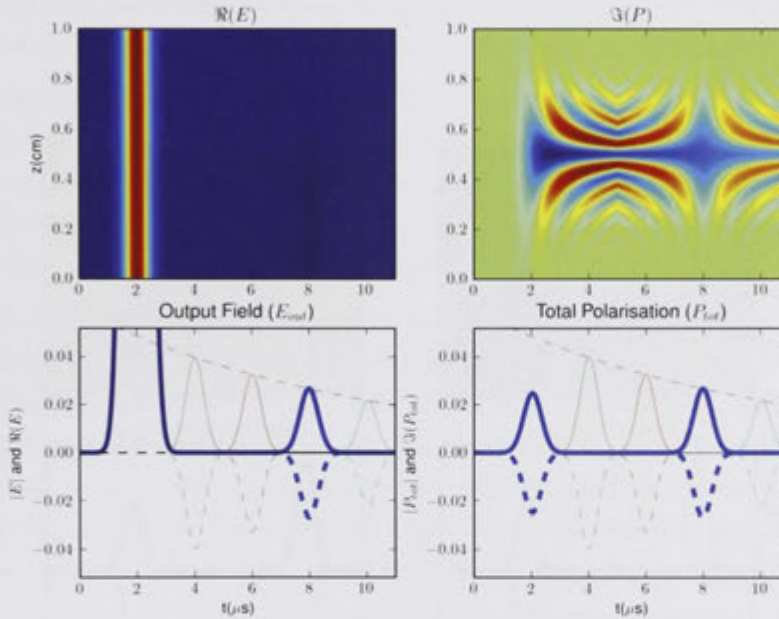


Figure 4.2: Simulation of a simple gradient echo, in the low OD regime. LHS: The real part of the E-field vs space (top) and time and the field at the input and output faces of the medium (bottom). RHS: The polarisation density in the medium (top) and the total polarisation summed over the entire medium (bottom). In the bottom panels, the solid lines show the absolute value of the field, while the dashed lines show the real part of the retrieved electric field (left) or the imaginary part of the medium polarisation (right). The thick lines correspond to the same simulation as the upper panels, while the fainter lines correspond to repeated simulations with different storage times. Lines represent the process repeated with different delays. The dotted lines shows the decay of an FID of the initial feature (an exponential decay for a Lorentzian). These are well described as a generic echo processes (section 1.3.11). Parameters:  $\beta_{\text{br}} = 0.2$  dB,  $\gamma = 0.1$  MHz,  $2Gz_0 = 4$  MHz,  $N(\Delta) = \delta(0)$ .

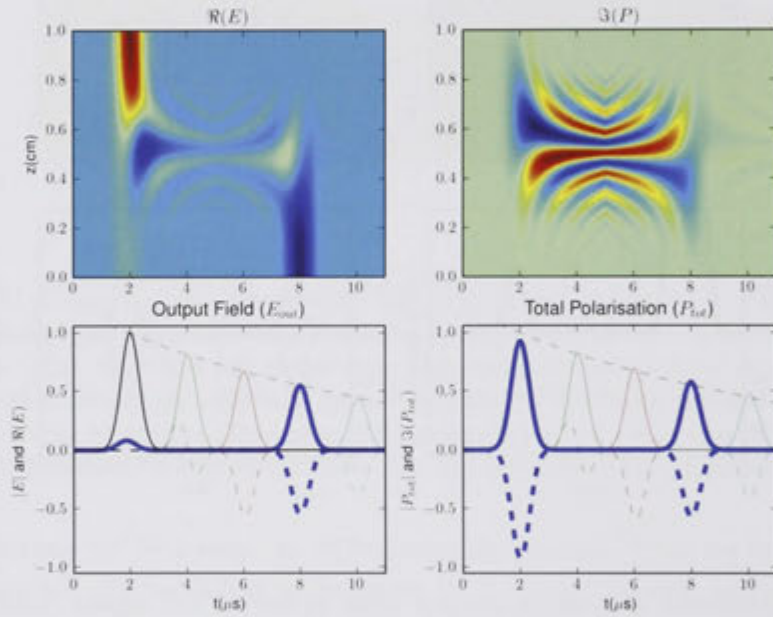


Figure 4.3: The same as the previous figure (see its caption), except now with a more absorbing medium- the optical depth of the broadened medium is now  $\approx 10$ dB. We obtain a much larger echo in this case, but with some phase distortion. This is because the input pulse is slowed by the medium and hasn't yet reached the centre when the field is switched. Parameters:  $\beta_{\text{br}} = 25$  dB,  $\gamma = 0.1$ MHz,  $2Gz_0 = 4$ MHz,  $N(\Delta) = \delta(0)$ .

in the mode of the retrieved pulse. This effect increases with optical depth, but decreases with storage time.

### 4.2.3 A Frequency Domain Memory

The gradient-echo may also be used to store modes separated in frequency. This may be beneficial in that the efficiency of each mode may vary less than the time domain case. It also lets us study the phase shift of the memory as a function of frequency.

We simulate the storage of temporally long, but spectrally narrow pulses in the optically thick regime (Fig. 4.4). We measure their phase variation by reference to a time-reversed copy of the input. Explicitly,  $\Delta\phi = \arg(E_{out}(t)) - \arg(E_{in}(-t))$ . We can compare this to the expected result for the limit of large broadening  $\beta \log(t_{end}/t_{start})$  where  $\beta$  is the optical depth of the broadened feature[58].

When the pulse bandwidth is well within the memory bandwidth, we find that the phase change is consistent with the result for infinite bandwidth (eq. (4.6)). However there is an additional constant shift which varies with centre frequency (Fig. 4.4d). This shift is not symmetric about zero detuning—the low frequency components, which are stored at the back end of the medium, are the strongest affected. This is because they experience the most delay during the input, and are thus ‘stored’ the furthest from resonance. In all cases, the most rapid phase shifts are at short times, appearing as frequency shifts.

The phase shifts introduced are deterministic, and thus reversible. A (potentially) simple way to do this is to cascade two memories. This works perfectly if they have infinite bandwidth, and the *same* detuning gradient. It has been previously stated[58, 136] that memories with *opposite* detuning gradients should be used, but this does not take into account that the output of the first memory is time-reversed, and thus the relationship of its spectral components is mirrored.

For finite bandwidth memories, cascading two memories will not completely remove the distortion, due to the frequency dependent component. The temporally changing component is still removed, as it may also be removed by applying an equivalent time-dependent but spatially-uniform detuning in the output step. If it is necessary to also remove the additional frequency dependent shift, a more complex, spatially varying modulation will be required, or a propagation through a medium with the correct dispersion profile.

It’s worth noting that even in the case of extreme distortion (well beyond that which has so far been realised in experiments) it may still not be necessary to alter the memory process. The phase shift is completely deterministic so that if we know the properties of the memory, we know the output mode based on the time since the field switch, and we might thus match our application to it. In particular, for applications based on photon number, the mode shape may be irrelevant.

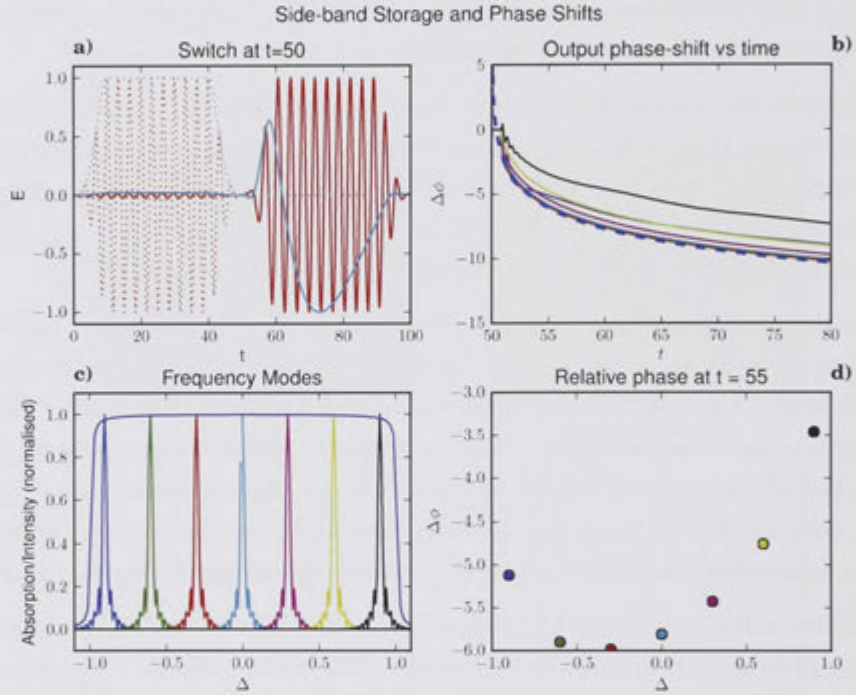


Figure 4.4: Simulated storage of pulses separated in frequency. a) The input pulses (dotted) are defined by a square pulse with a slow turn-on and turn-off to limit their bandwidth. The detuning is switched at  $t = 50$ . To avoid cluttering the figure, we only show two pulses with detuning  $\Delta = 0.3$  and  $0$ . b) The phase shift of the pulses vs time, calculated as  $\Delta\phi = \arg(E_{out}(t)) - \arg(E_{in}(-t))$ . The dotted line gives the phase shift for the large bandwidth case (eq. (4.6) with  $\beta_{br} = 7.5\text{dB}$  used in the simulation). The simulated lines are colour-coded with their spectra in panel c), where the memory bandwidth (dotted) is also given. d) The phase shift measured at  $t = 55$  as a function of frequency. Parameters:  $\beta_{br} = 33\text{ dB}$ ,  $\gamma = 0\text{MHz}$ ,  $2Gz_0 = 2\text{MHz}$ ,  $N(\Delta) = \delta(0)$ .

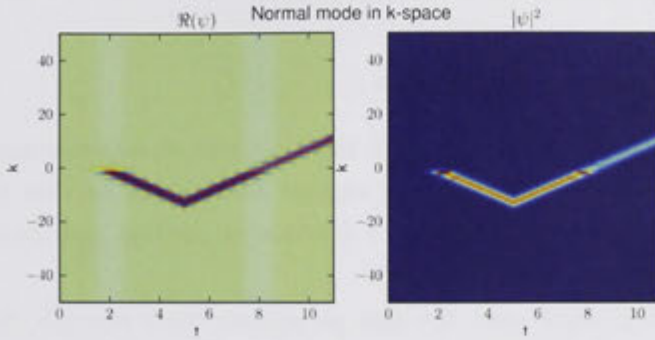


Figure 4.5: The time evolution of the normal mode,  $\Psi(k, t) = kE(k, t) + n\alpha(k, t)$  in  $k$ -space. This provides a nice visualisation of the evolution of coherence in the medium. Radiation is emitted when it crosses the  $k = 0$  axis- corresponding to the light moving at full speed to keep up with our co-moving reference frame.

#### 4.2.4 A Polariton in Time and $k$ -space

Another way to look at the process is to look at spatial oscillations. A normal mode then emerges naturally as a combination of field and atomic polarization. Following Hetet et al. [70], we Fourier transform equations 4.1 and 4.4 with respect to  $z$ . The normal modes are identified as  $\Psi(k, t) = kE(k, t) + n\alpha(k, t)$ , the equation of motion for which are:

$$\left(\frac{\partial}{\partial t} - \varepsilon(t)\frac{\partial}{\partial k} - \frac{ign}{k}\right)\Psi(k, t) = 0$$

$\Psi$  is a normal mode because all parts of it oscillate in phase. The evolution of this mode is governed by the detuning gradient, allowing a useful visualisation of how the coherence evolves in the optically thin regime. Figure 4.5 is similar to photon-echo focusing diagrams such as those used by Levenson[152]. This can be used to indicate which time in the past is currently being re-phased.

### 4.3 Gradient-Echo with Inhomogeneously Broadened Transitions

Lorentzian absorption features may be created either by homogeneous decay, or by a Lorentzian distribution of spectrally narrow absorbers. A gradient echo using a Lorentzian feature with either style of broadening gives similar output. Physically however, the situation is different as the spatial-detuning relationship means that light interacts with one side of the feature before the other. When we consider non-Lorentzian feature shapes, the output is quite different depending on the nature of the initial broadening.

We will take the limit of pure inhomogeneous broadening, such that any individual absorber has an infinitely narrow spectrum. The inhomogeneous line height is then pro-

portional to the spectral density of absorbers  $N(\Delta)$ .

### 4.3.1 Inhomogeneous Lorentzian

We repeat the simulations in section 4.2.2, this time with an inhomogeneously broadened feature. The inhomogeneous broadening suggest another way to view the process. We plot the polarisation density,  $N\Re(\alpha)$  as a function of position and detuning for several different times before and after the switch.

In the low optical depth case, the light propagates freely through the medium. The absorbers at any given  $\Delta$  are uniformly excited with depth. The coherence in the medium is a maximum at the peak of the input pulse (Fig. 4.6b). By the end of the pulse, the various parts of the medium have moved out of phase according to their detunings. After the field switch, the different regimes reverse their relative evolution, until the entire medium is again in phase and the echo is released. The small phase-shift accumulated across the initial line-width remains however, causing the efficiency to decay with increased storage time.

In the optically thick case, the dispersive part of the interaction becomes significant (as in Fig. 4.1 b and c). Just before the field is switched, the pulse is still interacting with absorbers toward the front end of the crystal, largely off-resonantly. The excitation is thus not symmetric about zero, and the echo is shifted to a higher frequency. The degree of the shift depends how far the pulse propagated, so that the frequency is time dependent.

We can also see that in the optically thick case, the light interacts mainly with atoms on the low frequency side of the initial feature. This has important consequences for non-Lorentzian features.

### 4.3.2 Inhomogeneous Top-hat

The propagation effects of the gradient echo are emphasised when dispersion is large- that is, when we have sharp changes in absorption. The extreme of such is a top-hat profile.

A top-hat type absorption profile is unlikely to be created by a homogeneous process (one might imagine an absorber randomly jumping between different detunings with equal probability). If such did occur, we would expect similar temporal decay behaviour as the Lorentzian case, that is given by the impulse response (Fourier transform) corresponding to a top hat shape- a sinc function, as opposed to the exponential for a Lorentzian<sup>2</sup>. An inhomogeneously broadened, approximately top-hat absorption profile is commonly created in hole-burning media.

We simulate such a profile for several optical depths, paying attention to the temporal decay. We find an ‘edge’ effect- as the optical depth increases, only one edge of the feature participates resonantly in the memory, and it is this which determines the decay time.

<sup>2</sup>Actually a sinc function modified to be causal- in this case just the positive side (see e.g. Warwick[45])

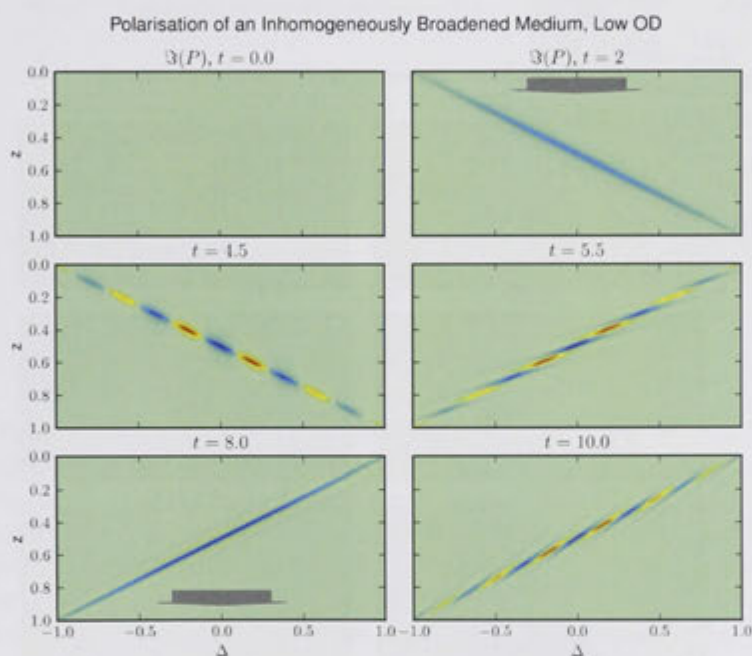


Figure 4.6: Snapshots of the polarisation of an inhomogeneously broadened medium, with low optical depth. Each panel plots  $\Im(P)$  vs detuning and depth at an instant in time, with red corresponding to positive and blue to negative values. A pulse with bandwidth  $\approx 1$  MHz is input at  $t = 2$ , and the field is switched at  $t = 5$ . This recalls the pulse at  $t = 8$ , although the medium remains excited as the coupling efficiency is low. Further detail in the text. Parameters:  $\beta_{\text{br}} = 0.02$  dB,  $2Gz_0 = 2$  MHz.  $\gamma = 0$  MHz,  $N(\Delta)$  is Lorentzian with  $\Gamma = 0.05$  (dephasing rate  $\gamma_{\text{inh}} = 2\pi\Gamma = 0.31$  MHz)

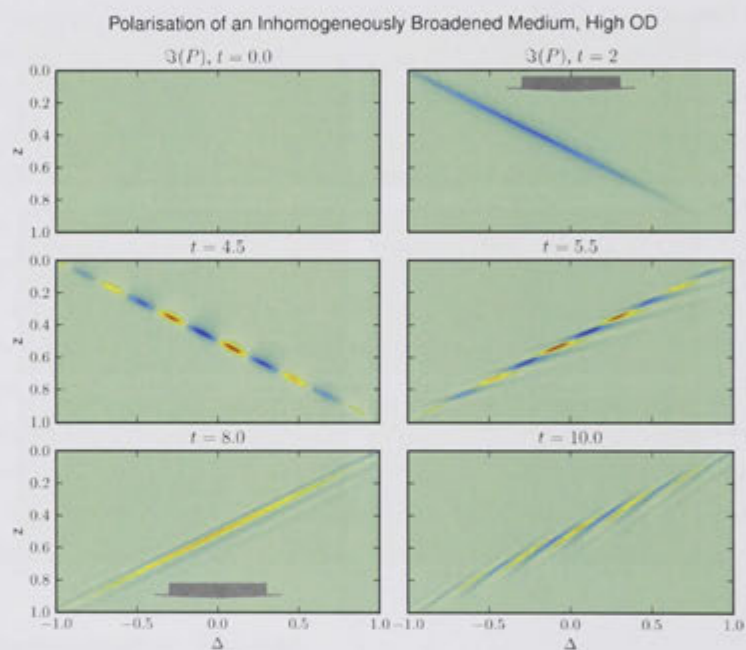


Figure 4.7: The same simulation as Fig. 4.6, but with higher optical depth. Propagation effects are evident in the second panel as the pulse excites the front of the medium first. Unlike the optically thin case, the pulse retrieval at  $t = 8$  de-excites the medium, leaving it close to the ground state. Parameters are the same as 4.6 except now  $\beta_{\text{br}} = 14\text{dB}$ .



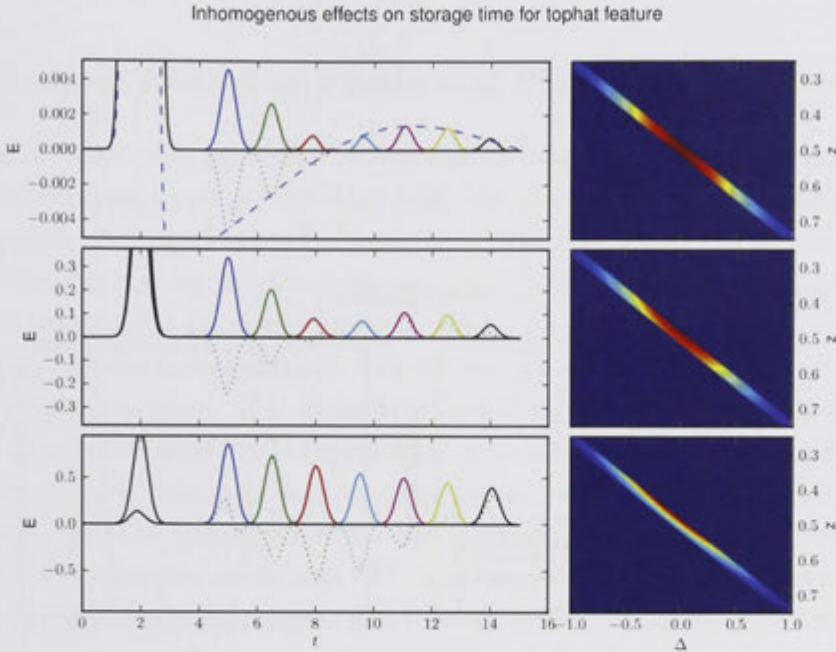


Figure 4.8: Echoes using a top-hat initial feature for increasing optical depth, and delays. Left: The continuous lines give the absolute value of  $E$ , the dotted lines show the real part. The dashed line in the upper panel gives the real envelope of the FID due to the initial feature when no field is applied. Right: The excitation in the medium at  $t = 4\mu\text{s}$ , before the detuning is switched. The smaller excitation width at higher ODs is responsible for the slower time decay. Parameters:  $\beta_{\text{br}}$  from top to bottom is 0.03 dB, 3 dB, 18 dB. All simulations:  $2Gz_0 = 3\text{MHz}$ .  $\gamma = 0\text{MHz}$ .  $N(\Delta)$  is a top-hat profile with width 150kHz.

The results are shown in Fig. 4.8. In the optically thin case, the echoes decay with the same dependence as the FID, i.e. sinc shaped. As the depth is increased, the decay time increases. This corresponds to a decreasing band of the initial feature (concentrated at higher detunings) being excited.

### 4.3.3 Hidden Transitions

A step further than inhomogeneous broadening is to have two completely resolved spectral lines, as may belong to two separate transitions. This situation is illustrated by a simulation with results in Fig. 4.9. We find that one is to some extent hidden by the other, in that in the optically thick regime, the only interaction with the hidden feature is off-resonant. This interaction causes a delay and a phase shift of the input pulse, but does not reduce its overall efficiency.

This is a significant result as it means that for efficient, wide bandwidth memory, it is only necessary to have a free spectral region on one side of the initial feature. In chapter 6 we discuss how this may be used to avoid bandwidth limitations of memories based on

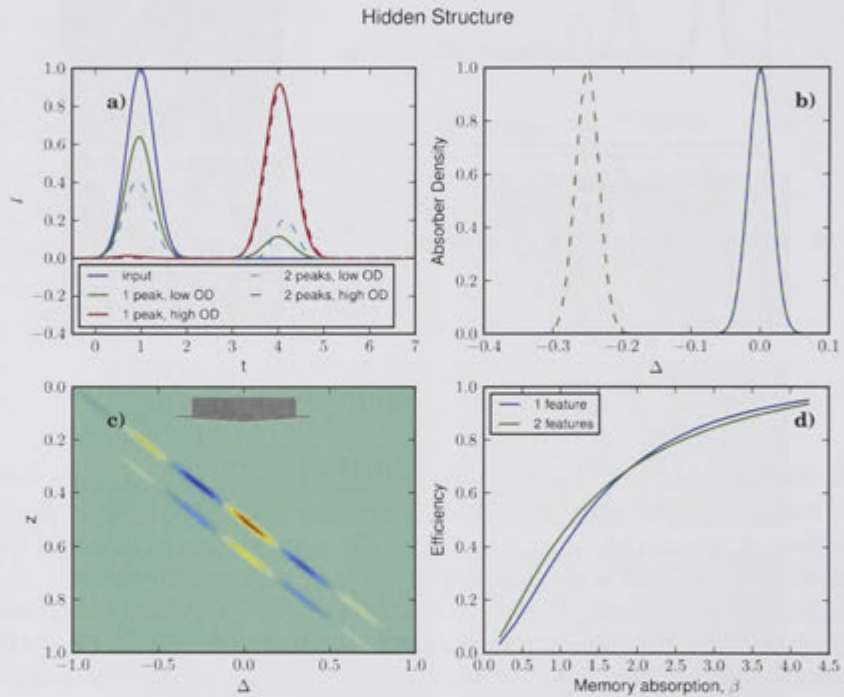


Figure 4.9: Performance of a gradient echo with the initial ensemble consisting of one or two absorption features. a) Input pulse and echoes with either one or two initial features compared, with low and high optical depths. b) The initial absorption features used in the simulation. The continuous line gives the absorption when only one feature is used. c) The polarisation of the medium immediately before the field switch, for the optically thick simulation in panel a), with two features. d) Echo efficiency vs memory optical depth for one and two features. Parameters:  $2Gz_0 = 2\text{MHz}$ ,  $\gamma = 0$ ,  $N(\Delta$  is given in b).  $\beta_{\text{br}}$  in a) is 2dB per feature (i.e. 2dB for 1 peak, 4dB for 2) for the low OD sims and 20B per feature for the high OD sims. For panel c),  $\beta_{\text{br}}$  is 7dB per feature.

hole-burning.

## 4.4 General Detuning Fields and Manipulating Light

So far we've varied only the initial feature and the shape of the input pulse. We now lift our assumptions about the detuning field, and considering variations both spatially and temporally. We make two main points here. The first is that the efficiency is quite robust to small deviations in an experiment. The second is that the detuning field offers a powerful window into the manipulation of light.

The gradient echo technique means that different spectral components are mapped onto different positions in space. This allows us to control the different spectral components as precisely as we can control the detuning field. If the detuning field is due to an external electric or magnetic field, this may be quite precise indeed. The method of control is quite clear in the low OD regime where the space-detuning mapping is direct. Things are slightly more complex in the high OD regime due to the additional time dependence, although full control remains possible. Some work in this direction has been reported using manipulations of switching speed and bandwidth as well as variable atom-light coupling in a Raman scheme[153, 71].

### 4.4.1 Shelving and Suppressing

In general it will be desirable to have immediate recall of an input pulse, without waiting for the full storage time again. This may be realised by a simple modification to suppress the atomic re-radiation and 'shelve' the evolution of the pulse in the memory. Related steps may be further used to increase storage time and increase the effective memory bandwidth. Further, the ability to suppress emission may be used for detailed light manipulations as discussed in the following section.

In the simplest solution, the evolution of the memory may be switched off by turning off the detuning field once the pulse is in the medium. The induced dephasing then stops, so that the pulse can be retrieved at shorter notice. In this configuration it is also easier to perform a coherence transfer to a shelving level for long term storage, as the spectral width that must be transferred can be very narrow.

This solution is not ideal however: a complication occurs when light is still propagating. This may be due to a large optical depth and a too-short storage time, or to an input that has a bandwidth larger than the memory.

This may be easily solved by a further step, at least theoretically. Once the time of the input pulse has passed, any light still exiting the medium is due to an FID, that is re-radiation by the atoms. To stop these radiating, one option is to induce large dephasing; instead of switching the detuning field straight off, we first increase it to many times its original strength for a short period. This both shortens and intensifies the FID. If the induced detuning gradient is strong enough though, the energy in the FID is reduced as




Figure 4.10: By turning off the detuning field, we cease the evolution of the memory to allow for a quick recall when required. In the high optical depth case, or when out-of-band components are present, some light may still be propagating through (green line). This may be removed by momentarily applying a large detuning field to completely dephase the medium (indicated by the total memory bandwidth,  $2Gz_0$ , lower panel). This must be reversed on recall, and may also recall components that would otherwise have been out of bandwidth. A similar result could be achieved by performing a rapid coherence transfer to a shelving level. ( $\beta_{br} = 22\text{dB}$  (when  $2Gz_0 = 1$ ),  $\gamma = 0$ ,  $N(\Delta) = \delta(0)$ ).

the effective optical depth of the medium is reduced. When we are ready to recall the light, we rapidly undo the extra dephasing with an equivalent gradient of opposite sign. Simulation results are in Fig. 4.10.

A more ideal solution to stop the re-radiation is to transfer the coherence away from the optical transition, to a free ground state. If a transfer may be performed with similar speed, the effect of suppressing the FID will be the same, with the added benefit that storage times may be many orders of magnitude longer.

Interestingly a step like this may in some cases be used to increase the efficiency of the memory, as more high frequency components may be stored.

#### 4.4.2 Field Switching Speed

We've so far been modeling the detuning gradient to change instantaneously, but this is unnecessary. A simulation is given in Fig. 4.11 to confirm that the memory efficiency is not strongly affected. There are propagation effects at high ODs as indicated by the slightly distorted pulse shapes, but these may in principle be accounted for.




Figure 4.11: In real experiments it is impossible to switch the detuning field instantly. We simulate a slow switching rate (indicated by the memory bandwidth,  $2Gz_0$ , lower panel) to confirm this has little effect. The retrieved intensity traces (upper panel) are matched in colour to the detuning fields (lower panel). ( $\beta_{\text{br}} = 20\text{dB}$ ,  $\gamma = 0$ )

### 4.4.3 Manipulating the Output Mode

By changing the detuning field for the output step, the pulse mode may be manipulated almost arbitrarily. This was first described by Buchler et al. [153]. As the light is mapped to the state of the atoms, it becomes easy to control.

For the low-OD limit or long storage times, the field components are mapped onto resonant atoms, so that there is a direct correspondence between manipulating the atoms and the field. For the high OD, low storage time regime, accurate manipulations must also take the delay into account. Doing this, one may further remove the phase-shift introduced.

One example is shown in Fig. 4.12, where the gradient at the output is changed to compress or lengthen the pulse. In this process, the energy in the output remains unchanged unless a large enough gradient reduces the optical depth sufficiently (thus effecting the coupling efficiency). Indeed, we've already used this for shelving in Fig. 4.10.

### 4.4.4 Repeated Retrieval to Increase Read Efficiency

In the low optical depth case, repeating the field switch many times can be used to extract more of the stored energy. In the optically thin case, each time the atoms are rephased, a fraction  $\eta_{\text{read}} = \sqrt{\eta}$  (The read efficiency, generally the square root of the single-switch total efficiency: see chapter 2) of the remaining energy is recalled. Every second switch,




Figure 4.12: Using asymmetrical input and output detuning fields, the mode can be manipulated. Here we simply change the gradient at the output to broaden or narrow the pulse. The total broadened width of the memory is given by  $2Gz_0$ , shown in the lower panel.  $\beta_{br}$  (during input) is 16dB.

the rephasing is occurring in the opposite direction. While the first, and all odd numbered switches recalls a time-reversed pulse, the second is forward in time. This principle was discussed by Longdell et al. [58]

We note that as the optical thickness is increased, the efficiency of every second echo may suffer from reabsorption if the initial feature is inhomogeneously broadened, due to the edge effect as in Fig. 4.8.

#### 4.4.5 Non-monotonic Fields and Multi-stage Memories

As a general principle, the detuning gradient does not have to be linear. The main requirement is that it be monotonic to avoid reabsorption[136]. A particularly dramatic example of non-monotonicity occurs naturally in experiments such as ours however, due to the symmetry of crystals and thus the presence of multiple dipole orientations. YSO, as used in our experiments, is one of these. Here we simulate a highly absorbing sample with two, oppositely shifting detuning gradients.

The results are shown in 4.13 and 4.14. The pulse is absorbed within the first half of the crystal, and on switching the field is reabsorbed by the second half. It is thus important to remove one of these components to obtain an efficient echo.

Interestingly, we can still get an efficient, but distorted echo in this case. We switch the field yet again to recall it from the back end of the crystal. The output is distorted mainly

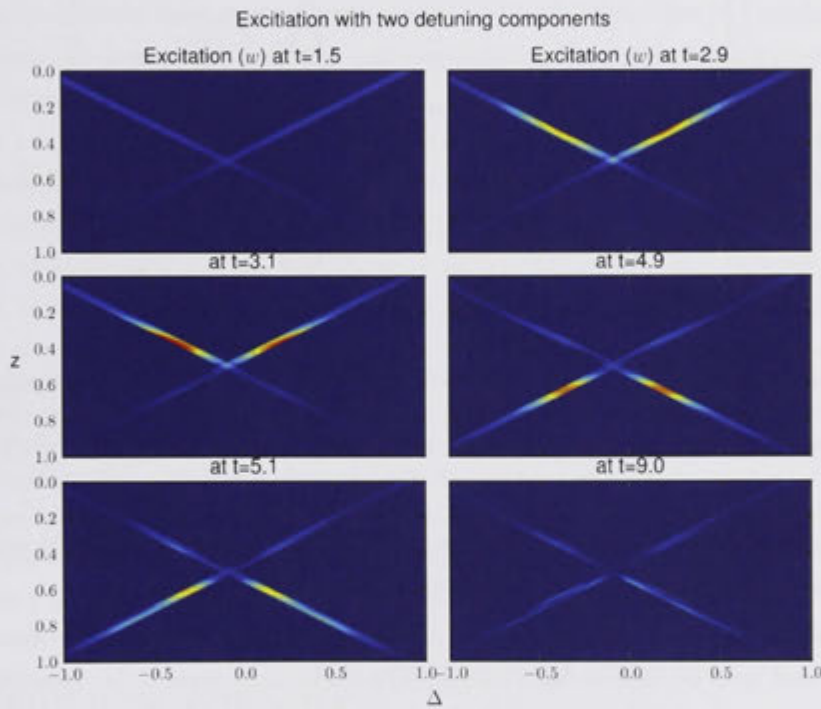


Figure 4.13: The excitation vs space and detuning for a medium with 2 sets of absorbers which shift oppositely under the detuning field. The snapshots are at the beginning, immediately before and after the first field switch, the same for the second field switch, and at the end. This clearly shows the progression of excitation from one end of the crystal to the other.  $\beta_{br} = 10\text{dB}$  for a single feature,  $20\text{dB}$  for both.  $\gamma = 0$ ,  $N(\Delta)$  is a top-hat, width  $\Gamma = 0.08\text{MHz}$ .




Figure 4.14: Intensity vs time for the simulation in Fig. 4.13. The detuning field is switched twice, but only on the second is a significant echo retrieved. The distortion is because the central frequency components of the pulse were not efficiently retrieved.

because the central frequencies were not recalled efficiently, but this may in principle be removed by detuning the pulse and/or using a different field pattern.

We note that having different shifting groups of absorbers offers an additional degree of control that may be used in future work. The shifts don't have to be opposite as here. If we have 2 centers with orthogonal dipole moments, a field gradient along each axis may be used to manipulate each center independently. Further degrees of freedom are given by combining other interactions, e.g. combining different Stark components and electric field with degenerate Zeeman levels and a magnetic field. In these cases, we may consider applying different field gradients to each group to affect their rephasing time, as well as constant shift to affect their overall phase.

Such techniques should find many applications. We might for example perform in-memory operations such as mode splitting and combining (along similar lines to, but more powerful than those discussed elsewhere[154, 155]), to split and combine frequency and temporal modes. This may be of use to create and to swap entanglement in repeaters, or just perform partial readouts. A particularly simple application is to temporarily shift two degenerate groups in order to give them a relative  $\pi$  phase shift. These will then interfere to suppress the formation of echoes, which might be used for pulse sequencing (below) or to perform quiet optical rephasing (section 6.2.2). This principle is the same as for Stark-modulated photon echoes which were originally used to study material properties([141]).



#### 4.4.6 Pulse Sequencing and Echo Suppression

It can be useful to re-order a series of pulses stored in the memory- in other words, give random access. This allows shuffling temporal modes for use in repeaters, as discussed in chapter 2.

As we have control of the spectral components of the pulse, it is a simple matter to rewind or fast forward through what is stored, simply by applying strong negative or positive gradients (or ideally a slight variation in the high OD case). To avoid affecting what's stored while we do this however, we need a way to suppress the output.

The principle has been demonstrated using a Raman transition in rubidium gas[71]. In this case, the rephasing is simply continued (using hyperfine broadening) while the coupling beam to create the Raman feature is turned off.

There are many ways to suppress the echo: anything which can remove the collective coherence or reduce the atom-light coupling without affecting the rephasing may be used. We have already mentioned two options: using a large gradient to reduce the coupling efficiency, or applying a  $\pi$  phase shift to cause destructive interference between degenerate levels. A further possibility is to use an additional broadening field, independent from the main one (e.g. using a field gradient along a different axis, or inducing a reversible phase shift optically using an off-axis beam), or combining the gradient with other echo techniques, for example an AFC.

We simulate this concept in Fig. 4.15, suppressing the output by using the simplest mechanism of a large broadening. We also use this to demonstrate retrieving one pulse without the time-reversal, although in the case of another transition hidden behind the first (as in Fig. 4.9), the time-forward retrieval suffers some reabsorption.

### 4.5 Improving Efficiency

So long as phase and mode distortions due to storage are not significant, or are correctable, the major consideration for any given application is the efficiency. Clearly we want to use conditions which maximise it. We restrict ourselves to the problem of storing a single mode for a brief time, but even here there is a large parameter space to consider. We might vary the pulse shape, the profile of the initial feature, or the detuning field both in time and space to obtain the best result. There are also many possible application situations, each with their own requirements and limitations. Thus a consideration of the general problem is beyond the present scope.

A relatively easy but useful problem is that of the optimum initial feature shape, with everything else fixed. This is a significant experimental parameter for our experiments, where spectral hole-burning may be used to choose it. We'll consider this first.

For experiments with a lot of freedom, the biggest parameter is probably the detuning field. In theory we could use it to perfectly match our medium to a pulse in time and bandwidth. This will depend on the pulse considered, and will need to account for




Figure 4.15: Suppressing echoes and changing the readout order. The detuning gradient provides an easy way to change the evolution of the ensemble, both forwards and backwards in time. If we can suppress the formation of echoes as we go through the time when they would normally emerge, we can change the pulse order. Here we use a convenient method of echo suppression, simply applying a large detuning field to move through the resonance rapidly and reduce the coupling efficiency. The detuning field is represented by the total broadening of the memory,  $2Gz_0$ , shown in the lower panel. The upper panel gives the relevant intensities. The spike at  $t = 6\mu\text{s}$  is a time-compressed version of the second pulse, very intense but containing little energy. This also lets us readout forwards in time. Doing this for the case of inhomogeneously broadened media, or a ‘hidden’ transition, the pulse encounters reabsorption. This is shown by the green line, where a second feature detuned by 0.2(MHz) but nominally ‘hidden’ is introduced (as in section 4.3.3).  $\beta_{\text{br}} = 10\text{dB}$  per feature,  $\gamma = 0$ .

dispersive effects in the optically thick case. Roughly speaking the optimal broadened feature will create a coupling efficiency weighted according to the pulse intensity in time and frequency.

We will instead look at a simpler problem, choosing a pulse shape for a given memory configuration. While perhaps less useful ultimately than matching the memory to the pulse, this is much easier to study experimentally as arbitrary pulse shapes are quite easy to generate (at least if we limit ourselves to coherent states). A brief consideration of this problem will also give us some insight into the more general optimisation problem.

#### 4.5.1 An Efficient Initial Feature

An easy situation to realise experimentally is the storage of Gaussian-shaped pulses with linear field gradients. In our memories based on hole-burning, we may decide the shape of our initial feature, but we are limited in the peak height by the medium used.

To make a comparison of the memory efficiency vs the feature shape, we simulate the storage of a  $1\mu\text{s}$  FWHM pulse, and switch the detuning after a  $1\mu\text{s}$ . This corresponds to a time when 99% of the pulse energy has entered the medium. We repeat this for Lorentzian, Gaussian, and top-hat feature shapes for increasing peak absorptions. The results are plotted in Fig. 4.16.

As we expect from previous discussion (section 4.3.2), the top-hat and Gaussian features continue to improve with increasing optical depth, while the Lorentzian asymptotes. This is because as the optical depth increases, the light interacts more with one side of the feature. As the Gaussian and top-hat shapes get spectrally sharper toward one side, their dephasing rates decrease. This emphasises the point that top-hat features are the best for high efficiency storage, at least when optical depths are high. We note that in real experiments, a homogeneous decay is also present. The high optical depth features cause slow-light effects, meaning the light remains in the memory for longer. This extra delay may incur greater loss due to the homogeneous decay, so that increasing the optical depth will reduce the storage efficiency. This is especially noticeable for the  $\Gamma = 0.2\text{MHz}$  Lorentzian simulations. We should note that this is only for exceptionally high optical depths however, and could also be removed by tailoring the output detuning field to avoid the delay.

#### 4.5.2 An Efficient Pulse Shape

The type of pulse that is best stored by a given memory is one which is temporally long enough to fit in the memory's bandwidth, but short enough to minimise decay over the memory's storage time. The fact that our memory is linear suggests that finding the best pulse is easy, but this is significantly complicated by the system's time dependence. This means that the frequency and time dependence cannot be separated. We might loosely think of this in one of two ways: when interacting for short times the atoms absorb over a larger bandwidth, or alternatively that out-of-band spectral components may be delayed

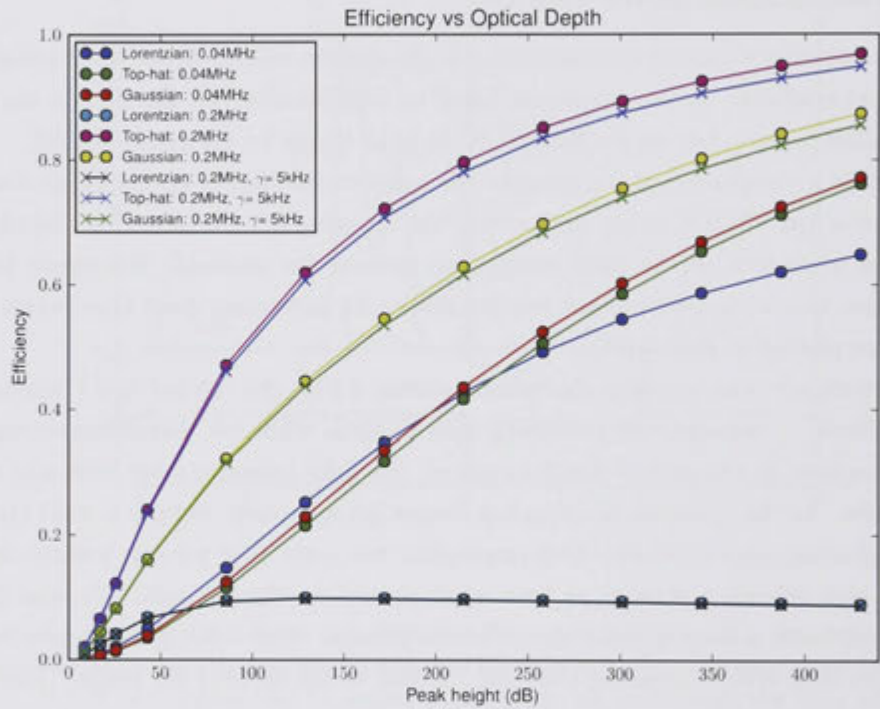


Figure 4.16: Efficiency of storage for Lorentzian, Gaussian, and top-hat initial feature profiles of various widths ( $\Gamma$  given in the legend), plotted against the height of the initial peak. These correspond to a Gaussian input pulse of FWHM  $1\mu\text{s}$ , and broadening of 2MHz switched  $1\mu\text{s}$  after the input peak. For the simulations marked by crosses, a homogeneous decay of  $\gamma = 5\text{kHz}$  was introduced.

by the memory, but not stored. The practical result is that the optimum pulse shape will be different for short storage times than for long times. Efforts have also previously been made to optimise the pulse shape for EIT[61].

This may also raises an issue with a definition of efficiency, as to whether we should include the out-of-band components as stored, seeing as we cannot control their delay in the same way as the in-band components. The issue is avoided (and illustrated) in section 4.4.1 by ‘shelving’ the out-of-band components. While this may not always be possible, it will be for many practical cases (when a large field, shelving level, or potential for causing interference between sub-levels is available), and will greatly simplify the present discussion.

To calculate the optimum pulse, we must know the properties of the memory. A time-invariant linear system may be characterised by a single measurement to obtain the impulse response. A gradient-echo is much more complicated, at least in the optically thick regime, due to the slow-light effects. The output is governed by the detuning field’s dependence on time and space. To characterise this system we would need to measure the impulse response at every point in time.

Rather than tackling this problem, we present a simple method to ‘measure’ the pulse, using our simulation or a real experiment. The premise is to use the output of the memory as a guess as to what it can store efficiently, then feed it back. We expect that the resultant pulse shape is optimal, although we don’t prove it.

We begin with a long pulse which also has high frequency components, such as square pulse. This is inefficient as much of the pulse is transmitted, or lost due to dephasing- but what is recalled represents both the bandwidth and temporal decay of the device. In the optically thin regime, it in fact represents all properties of the memory. If slow-light effects are significant however this is only partial information. We instead consider the recalled light as a guess at an improved pulse shape, and feed it back into the memory.

We take the measured output,  $E_{out}(t)$  and reverse it. As the gradient echo also introduces a phase shift, this means also taking the complex conjugate (we originally input a real pulse, thus the imaginary parts at the output are due to phase-shifts the storage). Specifically the signal we feed back is  $E_{in} = E_{out}^*(t_0 - t)$  where  $t < t_0$ , and zero elsewhere. After repeating this process many times, a steady state is reached where the output is the same shape as the input. An example simulation of a moderate optical depth memory with total broadening of  $2Gz_0 = 2$ , broadened optical depth of  $\beta_{br} = 7\text{dB}$ , and a damping rate of  $\gamma = 0.1$  is shown in Fig. 4.17.

The sequence rapidly converges, monotonically increasing in efficiency as it does. The final input pulse has an intensity profile that is identical (although time-reversed and larger) to the output it produces. As shown in Fig. 4.17(bottom left), the frequency spectrum of the pulse is also a trade-off. The pulse has much wider bandwidth tails than the memory’s efficiency spectrum, but these may be accommodated because of the short storage time. As we expect from previous discussion (4.2.1), the pulse spectrum is slightly

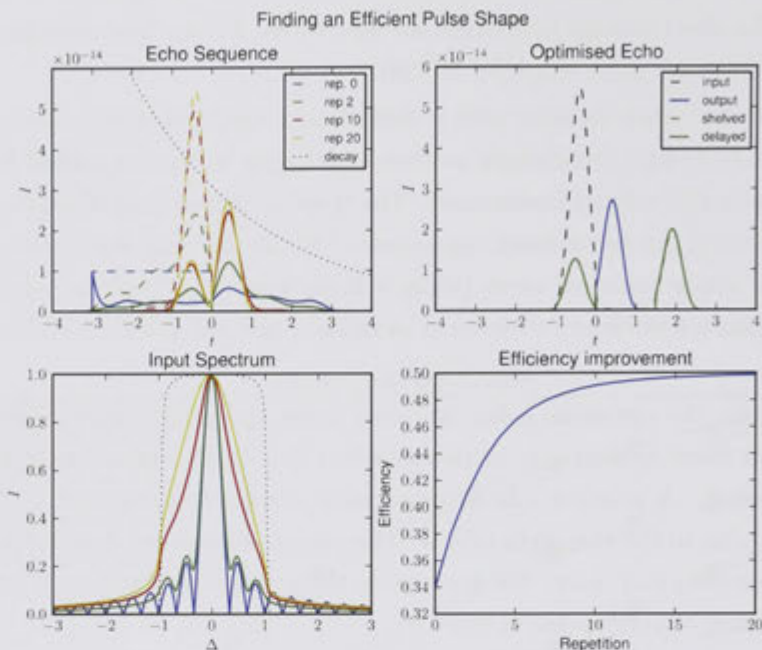


Figure 4.17: We estimate an efficient pulse shape for a memory by repeatedly taking the output and feeding back in a normalised, time-reversed, complex conjugate. We repeat this 50 times to arrive at a steady state. (top left) The input (dotted) and output (solid) every 10 repetitions. (bottom left) The spectrum of the input pulses. (top right) Storing the final pulse shape used (blue), with a delay (green), and with ‘shelving’ by first increasing the detuning field (dashed green). (bottom right) The efficiency of the storage as a function of repetition number. For all sims  $\beta_{br} = 6.5\text{dB}$ ,  $2Gz_0 = 2\text{MHz}$ ,  $\gamma = 0.1\text{MHz}$ ,  $N(\Delta) = \delta(0)$ .

asymmetrical due to the asymmetrical group-delay experienced. The final efficiency of  $\approx 0.5$  is limited largely by the available optical depth and dephasing rate of the medium.

To investigate the level of leakage caused by the high-frequency tails, we repeat the storage in Fig. 4.17(top right), but with a time delay (green line). As expected, there is some transmitted light, but it’s relatively very small. This is because out-of-band frequency components will tend to be out of phase with the in-band components. We also simulate the shelving case, where the broadening is instantaneously increased (by a factor of 20) to prevent the leakage. In an ideal experiment this step would be replaced by transferral to a long-lived ground state.

We haven’t proved that this is the most efficient pulse shape for this memory, but it certainly represents a reasonable guess. It is a weighted trade-off between the memory’s bandwidth and storage time. The same principles also apply to other techniques, such as an AFC or transverse CRIB. The lack of propagation effects in these techniques make them much simpler though, such that it is more practical to just calculate the optimal pulse shape directly.



---

# Realisation of an Efficient Gradient Echo

---

This chapter covers experimental work realising and characterising a high efficiency quantum memory based on the gradient echo technique. We use an absorption feature created by hole-burning in PrYSO, whose broadening is manipulated using an external electric field and the Stark effect. The bulk of the chapter deals with meeting the theoretical requirements for high efficiency: a high-contrast feature with optical thickness approaching 100dB, and a monotonic, approximately linear detuning field. The remainder deals with measuring efficiency and noise properties, and demonstrating some properties of the gradient-echo discussed in the last chapter.

We break down the chapter as follows.

In section 5.1, we cover technical details common to most experiments, and unnecessary to understand the physics. The main details are the experimental set-up (which is largely common to all experiments), and the solution of consistency/stability issues for both preparation and measurement.

Sections 5.2 → 5.4 focus on preparing the memory.

In section 5.2 make measurements of the hole-filling rate vs concentration and temperature in PrYSO. These measurements indicate that hole-filling is caused by ion-ion interactions, and that lower concentrations offer better hole-burning properties at the expense of needing a longer crystal. Based on these results, we choose a 2cm long PrYSO, doped to 0.005% with Pr<sup>3+</sup>.

Section 5.3 describes the creation and measurement of a suitable optical feature in this crystal using competing optical pumping and ‘cleaning’ steps. Propagation effects in hole-burning are avoided by preparing the crystal from the side, perpendicular to the direction of storage. Using these techniques, very sharp optical filters are created, up to 300dB peak height, with a roll-off of 10dB/kHz. A method to characterise these features by preparing and measuring only part at a time is also demonstrated.

Section 5.4 focuses on achieving a monotonic Stark shift. This is complicated in PrYSO by the crystal symmetry, with 4 site orientations and thus 4 electric dipole orientations. We select atoms in 2 of these orientations using hole-burning, but a very precise electric field alignment is required so that they Stark shift together. We make measurements of



the Stark shift and use finite element-based modelling to design electrodes which avoid sensitivity to misalignment. Including the crystal dielectric constant in this model and a fit to experiment data, we are also able to estimate the effects of crystal/laser misalignment on the observed shifts.

The final 3 sections (5.5→5.7) describe actual echo experiments.

In section 5.5 we demonstrate efficient gradient echoes based on the optimised preparation steps. The limitations of these experiments are described and supported with theoretical modelling.

In section 5.6 we make quantum-noise limited measurements to demonstrate that the memory is indeed acting at a ‘quantum’ level for the mode we are storing, and is further performing better than the no-cloning limit. We use results derived in chapter 2 for this purpose. A simple, (although slightly novel) tomography technique based on maximum-likelihood estimators is described to optimally use the available data.

Section 5.7 presents some other related experimental data on two phenomena predicted in the last chapter. These are the ‘edge’ effect and the use of two oppositely oriented Stark groups to create a two-stage memory.

## 5.1 Technical Considerations

Experiments described here have many common details, using essentially the same optical setup and similar measurement techniques. They all involve a crystal of PrYSO held at < 4K, prepared and measured using laser light. The details of this layout, and the precise control and measurement of the laser and sample were crucial to the experiments, but generally not necessary to understand the physics.

The equipment and optical layout are described before some details of measurement and preparation stability are focused on in depth. These considerations were especially important for the measurements of quantum noise in section 5.6, and for precise and consistent preparation of the sample when implementing a high-efficiency echo.

### 5.1.1 General Experiment Setup

The optical layout shown in Fig. 5.1 is used for the experiments described below, with only minor variations. The crystal to be studied is held in a liquid helium cryostat and maintained at approximately 4K in low pressure Helium gas, or immersed in super-fluid helium. Light was guided, gated, and modulated onto the sample crystal using AOMs driven by amplified digital synthesis sources. The rough configuration of the driving electronics is given in Fig. 5.2.

Three beam paths were used in the experiments, which we refer to according to their major role. These are the Probe, Burn-back, and LO paths. The sample was prepared by illuminating it using a combination of the burn-back and probe beams. The result was ‘probed’ by observing the transmission along the probe path. The probe beam’s power

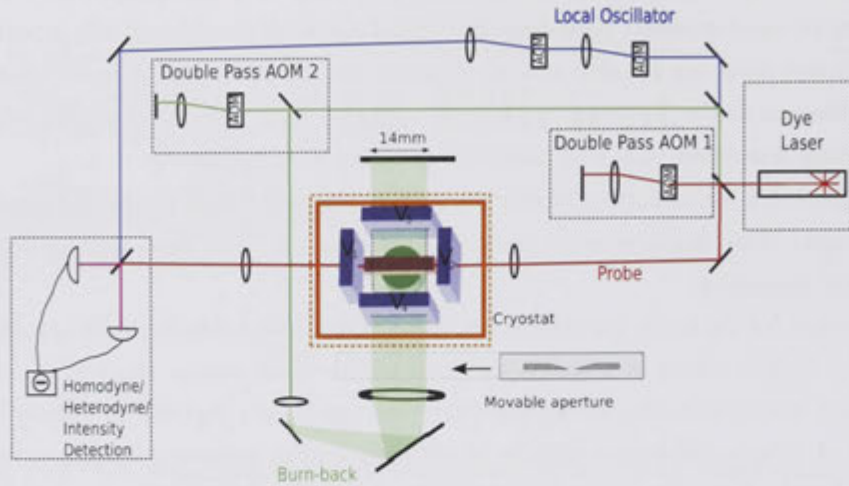


Figure 5.1: The optical table layout and all three beam paths used in the experiments. The probe beam is used in all experiments as it is the transmission/storage of this beam that is measured. The local oscillator is used when phase-sensitive detection is required, and the side-burning beam is used for avoiding propagation effects in the probe beam.

ranged from mW for preparation down to pW for quantum storage. This was focused with a 20cm lens to  $\approx 100\mu\text{m}$  at the centre of the crystal. The burn-back beam was a large, defocused beam of  $\approx 2\text{cm}$  width with the aim of illuminating the crystal uniformly. Its maximum power was typically a few mW. The local oscillator beam, when used, was mode-matched with the probe at the final beam-splitter for phase-sensitive detection. Its power was in the range 1.5-2mW. Intensity measurements, with no local oscillator present, used a 125MHz NewFocus detector, while homodyne and heterodyne measurements used matched homodyne detectors of a custom design, and a bandwidth of 5MHz.

Also employed for many experiments was an adjustable aperture on a translating stage, which allowed illumination and thus preparation of small sections of a large crystal at a time. Not shown are polarising optics to control the power in each path, maintain vertical polarisation at the crystal, and to match the probe and LO beam polarisations at the detector. Further, the ‘laser’ in these diagrams should be considered to consist also of various stabilising components (discussed below).

The base frequency used in the experiments, i.e. a ‘zero’ detuning, was 150MHz higher than the actual laser line unless otherwise stated, based on 2 passes of an AOM driven at 75MHz. This frequency was chosen to avoid aliasing, as many of the digital synthesis sources used 300MHz clocks.

The sample is mounted within a set of electrodes for creating electric fields. Two electrode geometries are described in 5.4, only one of which is shown here. Voltage was

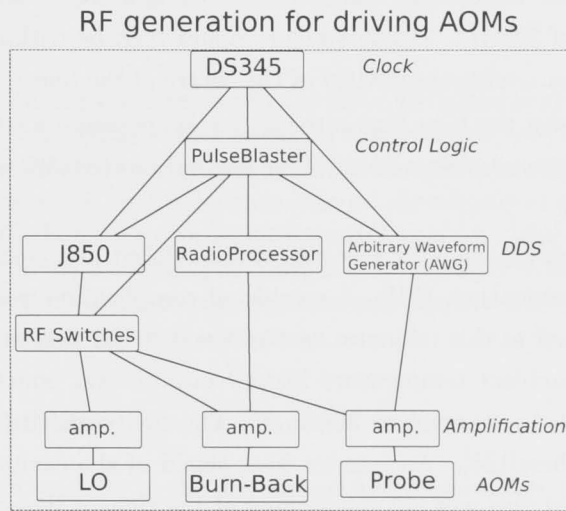


Figure 5.2: Rough schematic of the equipment for driving the AOMs. The entire system is synchronised to the clock of a Stanford Research Systems DS345 function generator. The major control logic is provided by a SpinCore PulseBlaster running at 100MHz. TTL signals from this are used to trigger the other components. The ‘J850’s are in-house built digital synthesis RF sources with a clock rate of 300MHz. Their main role is providing frequency-shift-keyed (scanning) RF. The RadioProcessor is also from SpinCore. It is used largely for its ability to shift rapidly between several frequencies, as required for burn-back. The arbitrary wave-form generator (AWG) is a Tektronix 510. It has a 1GHz clock with 10 bit amplitude resolution. It’s main purpose is to provide a quiet, flexible pulse sequence for storing/probing.

applied to the electrodes using grounded,  $50\Omega$  coaxial cable. These were supplied by a custom built device to switch each electrode rapidly between two power rails. The rails were up to  $\pm 50\text{V}$  and the switching speed  $150\text{ns}$  for 90%. Other voltage patterns could also be applied via a digitally controlled relay and a function generator.

### 5.1.2 Laser Stability

The light source was a highly stabilised dye-laser, described elsewhere[156]. This was tuned to  $605.977\text{nm}$ , corresponding to the centre of the PrYSO inhomogeneous line. In most experiments, this wavelength was confirmed using a highly sensitive wave-meter to a relative accuracy of  $30\text{MHz}$ . Measurements should then be within  $300\text{MHz}$  (the mode separation of the laser's reference cavity) of the centre of the line.

In terms of intensity, the laser fluctuated by several percent on a time-scale of seconds. A commercial intensity stabiliser reduced these fluctuations to  $0.5\%$ , at the cost of reducing the maximum power.

In terms of frequency, the laser was very stable, locked to a custom-built high finesse reference cavity to better than  $1\text{kHz}$ . A significant complication was that while the laser could be tightly locked to this reference cavity, it was found that at certain times of day (depending on the ambient temperature history over several hours), the cavity drifted relative to the sample by as much as  $30\text{kHz/s}^1$ . The cavity and its design properties are also described elsewhere[156]. An interior heat shield of the cavity itself was originally actively stabilised, but this was not functioning at the time of these experiments.

The drift was instead reduced by creating a stable temperature environment inside an insulating box surrounding the cavity. This simply used heating wire and a slow, digital PI loop to stabilise the cavity temperature near  $30^\circ\text{C}$ . With tweaking, this reduced the overall drift to at worst  $5\text{kHz}$  per second, although  $< 1\text{kHz}$  could be obtained for long periods (depending on temperature history). Reimplementing the active stabilisation of the cavity itself should reduce this much further.

Finally, some experiments, particularly noise measurements, required long stable periods for data acquisition on the orders of hours. To enable this, a slow pointing stabiliser was implemented. This device used two NewFocus quadrant detectors and 2 mirrors mounted on computer controllable mirror mounts (PicoMotor ) to correct the position and direction degrees of freedom of the beam exiting the dye laser. In this way, the sensitive alignment with the downstream interferometer could be maintained during long experiments. Noise on the beam due to the small, sudden corrections (PicoMotors being Piezo driven stick-slip actuators, moving in 'clicks') was detectable, but insignificant for our measurements.

The simple implementation is shown in Fig. 5.3.

---

<sup>1</sup> $>50\text{kHz/s}$  before its vacuum shield was re-vacuumed.




Figure 5.3: Simple pointing stabilisation for the laser, to reduce slow pointing drifts on the scale of minutes (and particularly after tweaking the laser cavity). The main requirement of the component placement was that mirror  $M_2$  is imaged onto detector  $D_1$ . This means that the alignment measured by  $D_1$  depends only on the alignment of  $M_1$ . A similar effort is made for  $M_1$  and  $D_2$ , but the requirement is much less strict. This is because the locking algorithm simply aligns  $D_1$  first using  $M_1$ .  $M_2$  is only adjusted (based on  $D_2$ ) when  $D_1$  is already aligned. For our locking to converge, we only require the signal at  $D_2$  to be *less* sensitive to  $M_1$ , not completely independent.

### 5.1.3 Sample Stability

The experiments were sensitive to the pressure on and temperature of the crystal, as these shift the resonant frequency of its ions. The extent of this was not quantified exactly, but variations of 10-100kHz were observed in some experiments. Most of this was apparently caused by fluctuations in the temperature and pressure of low-pressure helium in the sample space. This effect was reduced by adding ballast (a beer keg!) to the sample space and allowing super-fluid helium to cover the sample, thus connecting it to a larger thermal reservoir. Frequency stabilities on the order of the laser drift (1KHz/s) could then be obtained. This was measured by creating a feature and periodically probing its detuning over a 30 second period using a short pulse. Temporary deviations of 10kHz were still sometimes observed in these measurements, but the regular drift was removed. The cause of these temporary deviations is not clear. They appear to be irregular, causing broadening as well as shifting.

### 5.1.4 Interferometer Stability

The interferometer stability is a crucial factor for phase sensitive detection. The path length difference between the probe and local oscillator determines the phase at which the measurement is made. Typically this requires a stable path length difference, which may then be actively adjusted using e.g. a piezo-mounted mirror in the local oscillator path. The optical table used in these experiments was de-laminated during previous usage,

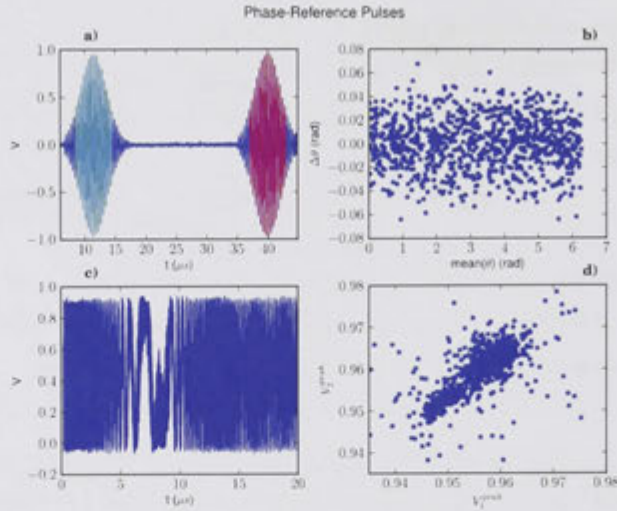


Figure 5.4: Interferometer phase measurement using reference pulses. a) By comparing before and after pulses we obtain a measure of the phase drift over the course of a typical measurement. b) Discrepancy in phase measurement giving the precision with which it works. c) A homodyne beat signal under the influence of the phase drift. d) Scatter diagram for the measured peak amplitude of each reference pulse. This indicates their use as amplitude references against the fluctuations of the laser.

causing such a scheme to be impossible. This is because vibrations were not damped and the path length drifted significantly. Attempts at preventing the interferometer drift by weighting and clamping the table were not effective.

The phase of the interferometer was instead simply allowed to drift, and measured at the required time. The measurement was done using detuned reference pulses (Fig. 5.4), one before and one after the time of interest. These were typically 2MHz detuned, where a transmissive spectral window could be created for that purpose. Such measurements were found to be accurate to 30mrad by looking at the difference in measured phase between the pulses (Fig. 5.4b). For the measurements shown here, the pulses are  $\pi/2$  out of phase, to avoid any common fitting bias. While not allowing actual control of the phase, the randomisation was sufficient for our experiments. It can be seen from the distribution of phase samples in Fig. 5.4d that the sampling is quite uniform, covering all phases equally.

In addition to the above it was necessary to improve the timing resolution offered by our recording oscilloscope. The device's (a 14-bit CleverScope) 100MHz clock was not able to be locked to the RF generation sources. As it samples 10ns intervals, trigger times may be out by as much as  $\pm 5\text{ns}$ , corresponding to an additional 50mrad at the 2MHz reference frequency. This mysterious piece of hardware, among its other features, is capable of missing occasional triggers by several clock cycles, resulting in greater deviations.

This trigger jitter was accounted for by recording a digital reference pattern at 1ns

resolution. This pattern was generated by the arbitrary wave-form generator also used to create the probe pulses. This allowed post-correcting the jitter with an accuracy of  $\pm 1.5\text{ns}$ , corresponding to  $\pm 15\text{mrad}$  at  $2\text{MHz}$ .

### 5.1.5 Detectors and Balancing

Highly sensitive measurements were important for many measurements, but especially crucial for measurements of quantum noise. The main issues for sensitive detection were to reduce classical amplitude noise on the local oscillator, and to calibrate and finely balance the detectors.

The primary source of noise on the LO was due to noise from the digital synthesis RF source driving its AOMs. By using an integer fraction of its clock speed ( $75\text{MHz}$  compared to  $300\text{MHz}$ ), and using passive low-pass filters at  $90\text{MHz}$ , most of the aliasing noise was removed. By also driving the source at full power and using passive attenuation to bring it down to the required level directly before the AOM, the effect of pick-up from other RF sources was reduced.

Balanced detection was used to subtract remaining noise. This was particularly important due to the relatively large amplitude noises on the dye laser and detectors at the low frequencies used here. In principle, arbitrarily good subtraction can be obtained provided the detectors are well enough matched.

Very good noise subtraction of  $40\text{dB}$  (in terms of measured intensity noise) could be achieved by tweaking the balancing beam-splitter, but typically this was not stable and soon faded to  $20\text{dB}$ . The consistency was improved by implementing pointing stabilisation, and by using digital subtraction (discussed in section 1.2.8), where the signal from both detectors (AC coupled) is recorded and the subtraction performed digitally. Digitally tweaking the balance to maintain was found to be good to maintain  $30\text{-}35\text{dB}$  of noise subtraction. The main limitation on this is the requirement on gain/signal size such that the scope could achieve its full  $14\text{bit}$  resolution above its electronic noise floor.

To calibrate the frequency response of the detectors, the relative phase and amplitude response spectra of the detectors was measured by comparing their response to short pulses to that of a high bandwidth detector. The amplitude part of this response is shown in Fig. 5.5a, with uncertainties of less than  $1\%$ .

Sensitive measurements were performed with the detectors just short of saturation. This corresponded to between  $700 - 900\mu\text{W}$  per detector, depending on the size of the beam on the photo-diode surface.

It remains that the biggest signal-to-noise improvement in these experiments would be better detectors with a lower electronic noise floor. As can be seen in Fig. 5.5, it is easily the largest source of noise at frequencies below about  $500\text{kHz}$ .

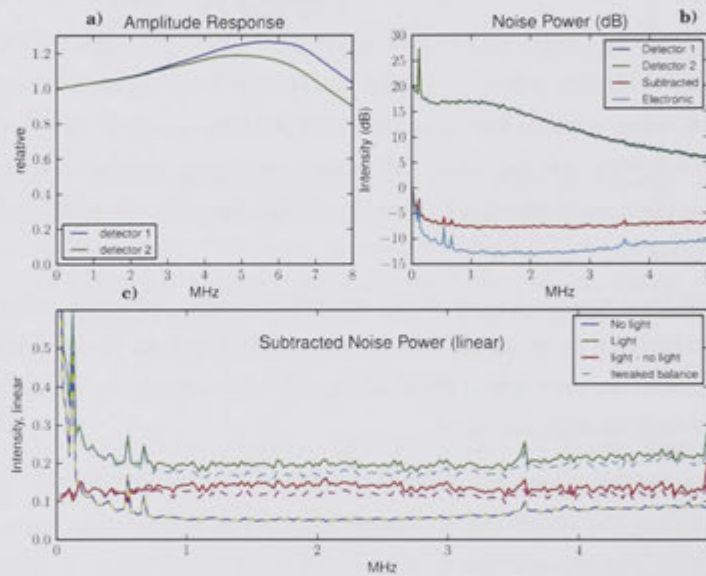


Figure 5.5: a) Homodyne detector response curves, as measured by comparing impulse response to a 125MHz NewFocus detector. b) Response-corrected spectra, and the electrical noise level. c) Digitally optimised subtraction, to minimise noise in the region 1MHz-2MHz. This included a time delay between channels of  $\approx 5$ ns, most of which was due to the oscilloscope's sampling.



## 5.2 Hole-filling in PrYSO

Having decided upon PrYSO as the medium, doping concentration and crystal size are two remaining important parameters for the sample. The larger the crystal and the higher the doping concentration, the greater the potential optical depth. On the flip side, higher OD means the trench will fill in quicker, requiring the experiment to be done faster, meaning more laser power and potentially power broadening. The hole-filling lifetime in PrYSO is however not well studied, and its main cause is not known. Further motivation to investigate the hole-filling was provided by some early echo experiments using a 2cm long PrYSO crystal doped at 0.02%. Spontaneous hole-filling in this crystal seemed to be the biggest obstacle in achieving highly efficient echoes.

Our main interest in the cause of filling is whether temperature or doping concentration play a significant role, as these are parameters we can change. Here we measure the hole-filling lifetime as function of temperature for two available crystals. These are both 4mm in length along the  $b$  axis with doping concentrations 0.005% and 0.02%.

The measurement method was to simply create a spectral trench by scanning an intense beam over 4MHz, wait a time between 1ms and 5 minutes, and measure the transmission of a CW beam tuned to the centre of the trench. The transmitted beam was measured directly on an intensity detector. This was repeated for different delay times, and also with a weak magnetic field due to a permanent magnet placed near the cryostat. Measurements vs temperature were made for fixed delays as the sample temperature was varied and logged. The samples were mounted one on top of the other, and both were housed inside an aluminium sheath to aid heat conduction. The temperature sensor was attached to this sheath.

The temperature data (Fig. 5.7b) indicate that temperature does not start to play a significant role until around 6K, and is thus not significant for our experiments. Above this temperature a strong dependence appears, roughly consistent with the  $T^7$  expected [157].

The concentration data however shows a strong dependence. To accurately determine the lifetime of a single ion, we need to account for the absorption at thermal equilibrium. The line shapes were given previously (Fig. 3.5). Care was taken to ensure our measurements were taken within 300MHz of line centre, with reference to a NewPort SuperCavity (for other measurements a much more accurate wave-meter was available). This indicates our measurements will be within 15% of the peak height in each case. Using the peak heights, we calculate the lifetime for a single spin (averaged over all the contributing transitions) as  $3000 \pm 1000$ s for the 0.005% crystal and  $700 \pm 150$ s for the 0.02% crystal. This indicates that the major cause of relaxation in these samples is the proximity between dopant  $\text{Pr}^{3+}$  ions, i.e. spin-spin coupling. This is supported by the finding that placing a small magnet near the sample could be used to influence the decay time. A naive model of resonant cross-relaxations is consistent with the factor of  $\sim 2$  increase in decay time that we observed—the field Zeeman splits the hyperfine levels in 2, reducing the number

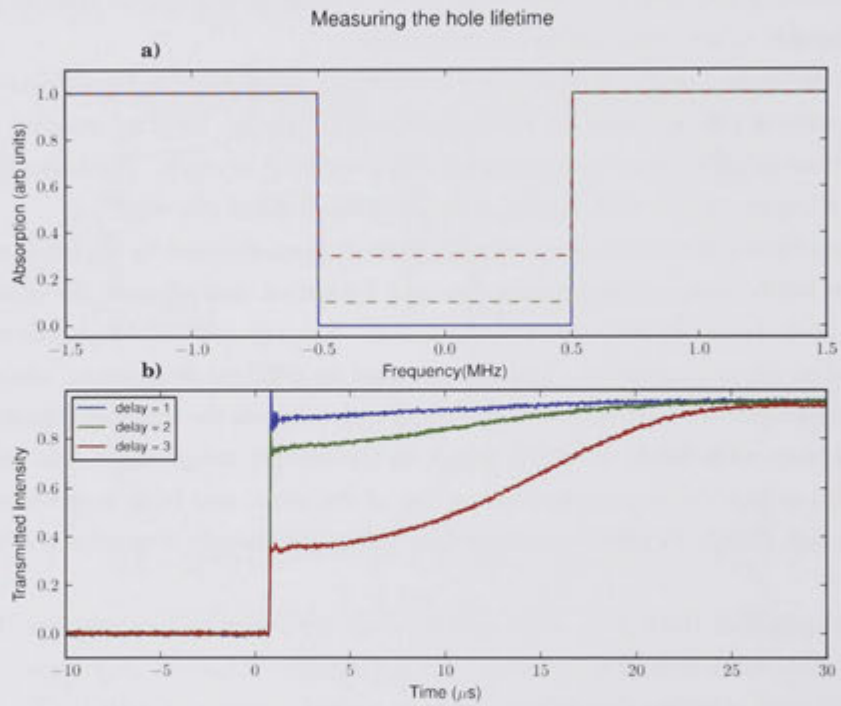


Figure 5.6: To measure the hole life-time, a trench was burned, some time was waited, then the laser (tuned to the trench centre) was turned on very weakly. a) A schematic of the hole filling. b) Some experimental traces using a 4mm, 0.02% PrYSO crystal. The average transmission during the period  $0.2\mu\text{s}$  to  $1.2\mu\text{s}$  after the laser was switched on was used to indicate the level to which the trench had filled.

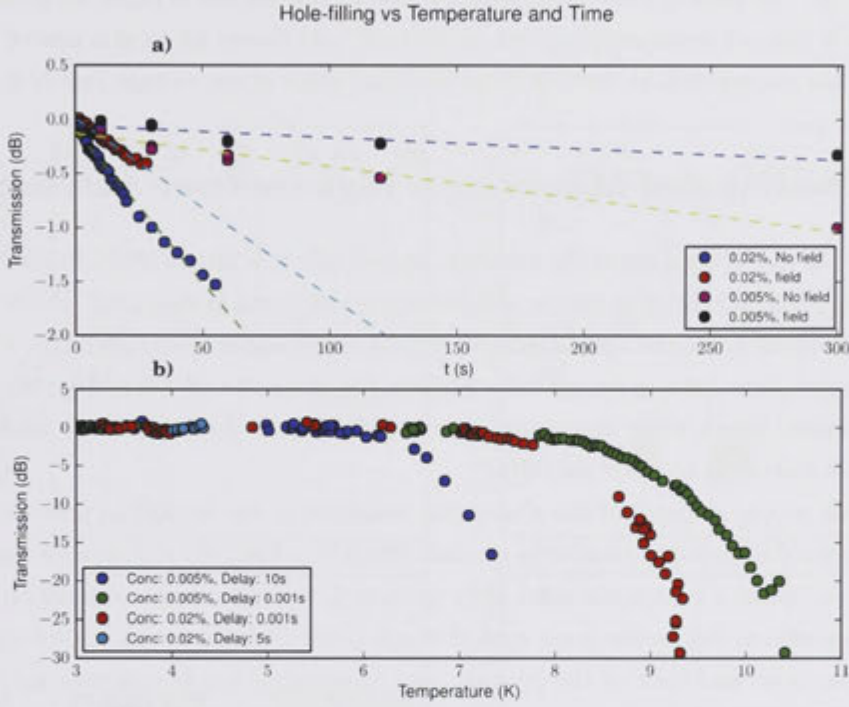


Figure 5.7: Measurement of hole-filling in PrYSO at a temperature of 3K (see also Fig. 5.6). a) Optical depth as a function of time, for different concentrations/magnetic fields. Least squares fits to the lines have gradients  $G = 28 \pm 2, 16 \pm 1, 3 \pm 0.5, 1 \pm 0.4 \text{ dB s}^{-1} \times 10^{-3}$ . Accounting for the peak absorptions in the crystals and the position within the inhomogeneous line, the expected peak absorptions are 54dB and 28dB. From these we calculate that an average ion will decay in roughly  $700 \pm 150\text{s}$ ,  $1500 \pm 300\text{s}$ ,  $3000 \pm 1000\text{s}$ ,  $6000 \pm 3000\text{s}$  for the respective cases. The magnetic fields were applied by placing a large permanent magnet outside the cryostat, about 20cm from the crystal and roughly perpendicular to the laser direction. The position and orientation of the magnet were chosen and varied to find a good increase in decay time. It is estimated to cause a  $13 \pm 5$  Gauss field at the crystal, simply by placing a Hall probe this distance from the magnet b) Optical depth as a function of temperature for a fixed delay. Samples were rejected if the temperature was changing faster than 0.05 degrees/s, or if it had changed more than 0.5K since the last measurement for the longer delays. The data indicate that temperature is not an issue for our experiments

of resonant nuclei seen by a given nucleus by 2.

### 5.2.1 Crystal Choice

Based on these measurements it was decided to use a 2cm, 0.005% crystal for further experiments. The peak of the inhomogeneous line of this crystal is expected to be 160dB (although it was not measured directly). A spectral hole should fill in at a rate of 1dB per minute. This corresponds to nuclear spins changing state at an average rate of 0.3mHz.

## 5.3 Creating and Measuring a High-contrast Anti-hole

We now focus on the creation of the anti-hole, specifically in a 2cm, 0.005% PrYSO crystal. Some techniques for achieving narrow anti-holes were discussed in chapter 3. Previous work has achieved both holes and anti-holes with widths on the order of kHz[158, 159, 124, 123]. However these have been quite optically shallow (on the order of a few dB). We aim for  $\approx 100$ dB optical depth while maintaining similar sharpness. This requires a filter with a much faster roll-off in terms of dB/kHz.

We wish to use as much of the absorption available in our crystal as possible. While the maximum theoretically possible is around 480dB ( $3 \times 160$ ), we will need to sacrifice at least half to obtain a monotonic stark shift (section 5.4). Other major considerations are propagation effects due to the large optical depth (such that the created feature changes between the front and back of the crystal), and broadening due to inhomogeneity in the ground state transitions.

The sequence shown schematically in Fig. 5.8 was settled on. It uses multiple burn-back frequencies to utilise all possible transitions and a competing process to create the feature and sharpen its edges progressively. Most of the preparation is done from the side of the crystal to avoid propagation effects. The sequence robustly approaches the theoretical limits of height and sharpness given enough time— and given perhaps the more stringent requirement of an intense, frequency-stable laser. This requirement of an exceptionally frequency stable laser is due to the slow preparation. Each step is described in detail in the following sections.

A final difficulty associated with the optically thick features was that of characterisation: to resolve the peak height of a 100dB thick feature using a direct measurement, a signal to noise ration of around  $10^{10}$  would be required. In section 5.3.6 this difficulty is avoided by preparing and measuring only part of the feature at a time.

### 5.3.1 Making a Transmissive Spectral Trench

The wide, deep spectral trench is made, simply by scanning the laser at high intensity over the spectral region to be emptied. The maximum spectral width is limited in PrYSO by hyperfine structure to around 18.5MHz. (The total range any ion can be moved is

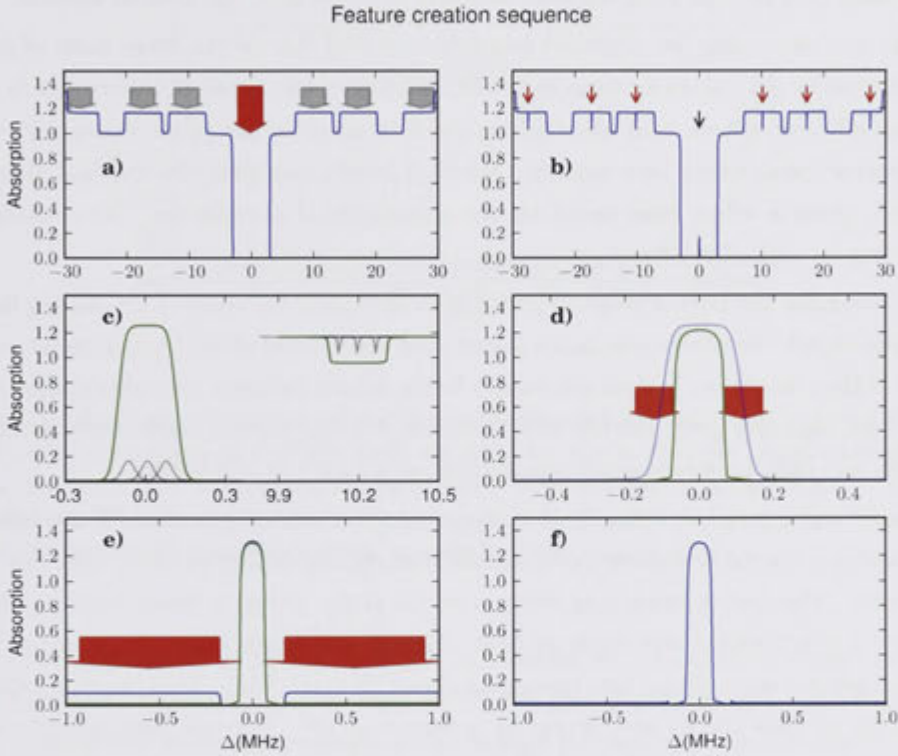


Figure 5.8: The steps used in creating a highly absorbing, top-hat shaped anti-hole. Red arrows denote applied laser light, while black arrows represent subsequent spontaneous decay. Propagation sensitive steps are applied from the side of the crystal, along the ‘burn-back’ path (Fig. 5.11). a) A trench is created by scanning the laser around zero detuning. b) Burn-back is applied to 6 frequencies, which are linear combinations of the hyperfine splittings. c) While this burn-back is in progress, a small, uniform electric field is applied to Stark shift the sample, creating a roughly square feature. d) Between burn-back periods, chirped ‘clean-up’ beams are applied to narrow the edge of the feature. e) Following the preparation sequence, extra ‘clean-up’ is applied along the probe path to keep the back-ground level low. f) The final feature.

17.3+10.2=27.5 MHz. But the worst case ions will also have a transition displaced oppositely by the total excited state splitting, of 9MHz. Thus they can only be moved 27.5-9=18.5MHz) (section 3.3.7). However we have a more stringent limitation in this case, as we wish to have a feature in the middle of this trench consisting of all possible transitions. Some of these ions will also absorb at a separations of the smallest excited state splitting, i.e.  $\pm 4.4$ MHz. In our experiments the trench is made by scanning the probe beam over several MHz with an intensity of a few mW, for several seconds.

This procedure may be expected to work very well due to the large ratio of ground-state lifetime to optical decay time in PrYSO. However, the depth of this trench is crucial to attain efficient echoes, and we will also use it to monitor the probe beam with optical and detector losses taken into account. We thus briefly calculate the residual absorption expected, given a worst case based on our experimental parameters. We consider ions interacting on and off resonantly.

To determine the level of residual population we equate the rates of ions falling into and out of the trench. We drive ions much faster than the excited state decay rate ( $\Omega = 1$ MHz vs  $\gamma = 1$ kHz.) Each ion is thus constantly being driven between the ground and excited states, such that any given ion has a 50% chance of being excited state at any given time. Each ion will then undergo spontaneous decay at a rate  $\gamma/2 = 0.5$ kHz.

Whether a decaying ion falls out of or back into the trench is governed by the branching ratios between the excited state and ground states corresponding to being out of or within the trench. The lowest branching ratio is  $\approx 1\%$  [137], giving a lower limit on the rate at which any ion group will decay of 5Hz. This is still much faster, by a factor of  $10^4$ , than the ground state decay rate measured above ( $3 \times 10^{-4}$ Hz). Thus we can reduce the population of even the slowest group by a factor of  $10^4$ . For our 2cm crystal we then expect an upper limit of  $10^{-2}$ dB (0.1%) to remain due to resonant absorption. This level will be reached at least as fast as the slowest decay rate (5Hz), and thus burning for a few seconds should be enough.

To calculate off-resonant absorption, we assume there is negligible population resonant within the trench, but outside the pumped window it proceeds as a step-function to the maximum. Following section 3.3.1, the absorption spectrum of this population will be

$$A = A_0 \frac{1}{\pi} \left( \tan^{-1} \left( \frac{-\omega - \Delta_0/2}{\gamma} \right) + \tan^{-1} \left( \frac{\omega - \Delta_0/2}{\gamma} \right) \right),$$

where  $A_0$  corresponds to the crystal absorption at equilibrium (160dB) and  $\Delta_0$  is the width of the trench. This is plotted in Fig. 5.9.

Moving in from the trench edges by 1MHz, the absorption factor reduces to 0.5%, or  $< 0.1$ dB. Thus if we create an 8MHz trench, the centre 6 MHz should be left with  $< 1\%$  absorption. This is less than the sensitivity of our measurements, although it may have implications for using larger crystals.

A measurement of the trench depth achieved in practise was also made. The probe




Figure 5.9: A rough calculation of the expected absorption remaining in a trench due to off-resonant absorption in PrYSO. Here we consider the worst-case population profile of a step function. The absorption in the centre 2MHz is a factor of at least  $1 \times 10^{-4}$  less than the equilibrium absorption, even for a relatively narrow 4MHz trench. If the peak absorption is 100dB, this corresponds to only 0.05dB, that is  $< 1\%$ .

beam with intensity  $\approx 3\text{mW}$  was scanned over 4MHz for 3 seconds to make the trench, then the average transmission of 10 weak ( $\approx 10\mu\text{W}$ ) pulses of duration  $1\mu\text{s}$  were measured. By comparing this to the transmission with the laser far detuned from the absorption line, the residual absorption was determined to be  $0 \pm 2\%$ .

### 5.3.2 Using All Possible Transitions

For many experiments using an anti-hole, there is a requirement that it consists of a single transition, for instance so that all ions react with the same Rabi-frequency. In our experiment we don't care about Rabi frequency, only requiring the largest optical depth possible. All transitions in PrYSO occur to some extent, and have the same Stark coefficient and coherence properties. We aim to create a strongly absorbing anti-hole consisting of all transitions.

This is achieved by burning-back at combinations of the ground state separations, shown in Fig. 5.10. As there are three equally populated ground states, the maximum absorption that can be reached in this way is 3 times the initial absorption (3.3.7). For our chosen crystal, this is 480dB.




Figure 5.10: The energy levels that are addressed to maximise the absorption of the anti-hole. Ions in PrYSO may be resonant at zero detuning on any one of the 9 possible transitions. The ions can be pumped into the correct ground state using only the 6 frequencies shown. A more optimal scheme would account for the differing branching ratios and oscillator strengths.

### 5.3.3 Avoiding Propagation Effects by Preparing from the Side

Hole-burning experiments typically use a single beam path for all preparation as well as measurement. As we wish to create an anti-hole in excess of 100dB however, such a practise is severely limited by propagation effects. The intensity of resonant light is already reduced by a factor of 10 after travelling through only 10th of the crystal.

Using coherent optical transients such as  $\pi$ -pulses may reduce these propagation effects to some degree, but have not been demonstrated with the optical depths we require here. What's needed is a way to effectively reduce the optical thickness of the ions for the preparation step, allowing all ions to be addressed, while maintaining high optical depth for the experiment. A simple solution is to do most of the preparation with light incident from the side of the crystal instead.

This works especially well as it not only reduces the physical thickness of crystal the burn-back light must penetrate through, but in PrYSO this axis is also less strongly absorbing. A 0.005% doped crystal was measured to absorb only  $0.2 \pm 0.1$  as strongly along this axis, that is  $2 \pm 1\text{dB/mm}$ . As the crystal used here is only 4mm wide, the burn-back beam only has to burn through at most 4dB to reach the centre, as opposed to  $\sim 100\text{dB}$  from the front.

In fact, if the experiment is repeated, the absorbing width is reduced to that of the probe beam. This is explained in Fig. 5.11. When a feature is made, the burn-back beam






Figure 5.11: By burning from the side, the optical thickness that the burn-back beam must propagate through becomes negligible. a) A fresh crystal in thermal equilibrium. b) The burn-back beam is applied from the side, eventually making the crystal transparent at the frequencies it addresses. c) An empty spectral trench is burned along the narrow probe path, filling in the holes previously made at the burn-back frequencies. d) The majority of the crystal remains transparent to the burn-back beam, with the exception of a few hundred micron where the trench was burned. e) Prepared crystal, with green showing transmission at the burn-back frequencies, red at the probe frequency.

creates holes at the frequencies it addresses. When the experiment is repeated (before the holes have had a chance to fill, i.e. within minutes), these holes are only filled in along the probe path.

This side burn-back (Fig. 5.1) was initially implemented using a broad beam of width 2cm, and a cylindrical lens to focus down to a narrow band  $\sim 1$ mm thick. The cylindrical lens was later removed. This avoided the problem of maintaining overlap with the probe beam, and resulted in an observed Rabi frequency of several kHz (i.e. just greater than the spontaneous decay rate) at maximum power. The intensity actually used in this path was reduced until power broadening wasn't observed.

Another potential method for reducing propagation effects was investigated briefly. This is to use a preparation beam along the probe-path, but with a rotated polarisation. For PrYSO this reduces the absorption seen by a factor of 10, which may be enough to allow preparation along the probe path depending on other factors such as birefringence. In such a case, a Pockels cell at  $45^\circ$  or similar would be required to switch quickly between polarizations for preparation and measurements steps.

### 5.3.4 Robust Preparation by Competing Processes

We ideally wish to make a top-hat feature. Using simple, narrow-band burn-back gives us a feature width limited by inhomogeneous broadening in the hyperfine levels— for PrYSO this is  $\Gamma_{\text{inh}}^{\text{hf}} \approx 30$ kHz. The peak optical depth is also reduced. Using spectrally

narrow burn-back to burn a deep hole of width  $\Gamma$ , the peak absorption of the corresponding anti-hole is reduced by a factor  $\Gamma/\Gamma_{\text{inh}}^{\text{hf}}$ .

To obtain the maximum OD, one solution is to scan the burn-back field over a larger bandwidth. For example, scanning over 100kHz will result in a  $\sim 85\text{kHz}$  wide approximately top-hat anti-hole, but with soft edges rolling off at  $\sim 30\text{kHz}$ . To gain the optimum sharpness, these edges may then be sharpened using another scanned field.

To implement this in our current experiment, appropriate scanning RF sources to drive and sweep the AOMs near 8 different frequencies would be needed— 1 for each of the 6 burn-back frequencies, and 1 to sharpen each side of the final feature. These were not experimentally available.

As an alternative, we used spectrally narrow beams, and scanned the detunings of the ions using the Stark effect. This was achieved by applying a small, spatially uniform but time varying electric field during the burn-back for a weak, dithered, 100kHz Stark shift ( $\approx 1\text{Vpp}$  at 1kHz applied to the electrodes, see next section).

Using a dithered electric field, no scanned or pulse-shaped RF fields are necessary. Practically however, the spatial uniformity of the electric field could not be guaranteed (see next section, 5.4). The frequency precision of the sharpening step was particularly important, so it was performed using the available RF sources to scan the laser. The burn-back steps were performed using the dithered Stark shift.

Now we could achieve the feature creation in only two steps: burn-back a tall feature, then sharpen. To achieve the optimum, it would be necessary to carefully choose the timing and light intensity of each step. In the present circumstances where we have a very spatially non-uniform burn-back beam and transitions with quite different oscillator strengths, this is nearly impossible. The solution we use is to treat the steps as competing processes as described in section 3.3.5, repeating them in short intervals many times. A further optimisation would be to match the relative intensity at each burn-back frequencies according to the transition strength, but wasn't attempted here.

The long-time limit for this process in the weak burning regime can be calculated as in section 3.3.5. We consider a single other ground state for simplicity, noting that the generalisation to more states is qualitatively the same.

For convenience, let  $T(\Delta_1, \Delta_2, \omega) = \frac{1}{\pi}(\arctan(\frac{\omega-\Delta_1}{\Gamma}) - \arctan(\frac{\omega-\Delta_2}{\Gamma}))$ . Setting the clean-up step to scan from  $\Delta = -200\text{kHz} \rightarrow -75\text{kHz}$  and the same for +ve detuning, the rate at which the unwanted ions are excited is proportional to

$$R_{\text{out}}(\omega) = \Omega_1^2(T(-0.2, -0.075, \omega) + T(0.075, 0.2, \omega))$$

with  $\Omega_1$  determined by the Rabi frequency on the relevant transition.

Conversely, we set the burn-back to scan over 150kHz, such that the rate at which ions are pumped back in is

$$R_{\text{in}}(\omega) = \Omega_2 T(-0.075, 0.075, \omega)$$

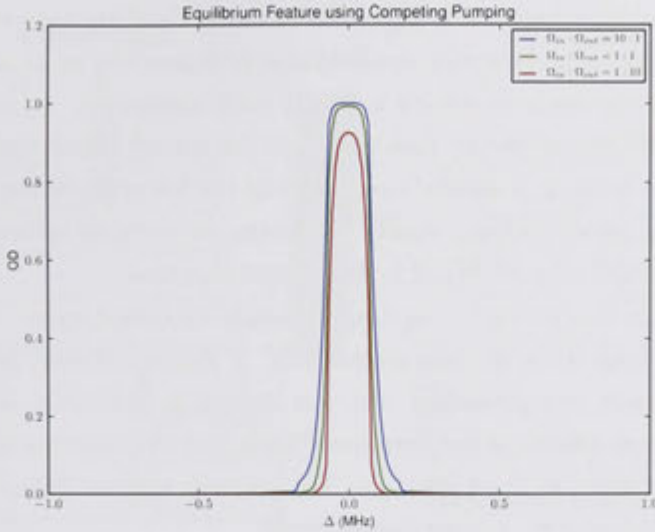


Figure 5.12: The resultant feature with competitive burning away and burning-in, over 3 orders of magnitude of variation in Rabi frequency. A broadband background clean-up is also applied.

These parameters are chosen to be similar to those used in experiments. Using the above expressions, and setting  $\Gamma_{hf}^{inh} = 30\text{kHz}$ , eq. (3.2) is plotted for three different ratios of Rabi-frequency in Fig. 5.12. If we were to include other transition pairs, the net result would be qualitatively the same—the features created on each transition would simply sum. If the branching ratios were determined solely by oscillator strength,  $R_{in}$  and  $R_{out}$  will be governed purely by light intensity such that the profiles would all be identically shaped. In practice there will be other decay paths involved however, which should be taken into account for a thorough calculation.

### 5.3.5 Cleaning up the Background

By sequentially applying the burn-back and sharpening steps, the feature height may be gradually increased while remaining sharp-edged. After many repetitions however, background absorption was found to accumulate outside the range of sharpening beams. This was mostly due not to spontaneous decay, but to frequency noise on the burn-back laser. An additional ‘clean-up’ step was added to reduce this background. This was applied along the probe-beam path throughout the preparation sequence, and detuned to just overlap in frequency with the scan range used to sharpen the feature. The motivation of this step was mainly practical rather than necessary: if the laser were to remain stable, adding this clean-up theoretically allows an arbitrarily long time to perform the other preparation steps.

### 5.3.6 Estimating an Optically Thick Feature

To understand our experiment, it is important to characterise the type of feature created. A standard transmission/absorption measurement is impossible, as an accuracy of 1 part in  $10^{10}$  would be necessary to resolve a 100dB peak absorption. In addition, a simple transmission measurement reveals nothing about the spatial distribution of the feature, only revealing the total. It is natural then to study the feature creation spatially. This is done easily in the present scheme simply by placing an aperture in the burn-back path, such that only a small slice of crystal is illuminated at a time.

A tall anti-hole was created using the previously described steps. This was further selected for a uniform stark shift(see section 5.4). A single (balanced-heterodyne) transmission measurement was performed, but was limited in sensitivity to about 30dB. A 0.5mm aperture was placed in the burn-back beam and the experiment repeated as the aperture was translated at 2mm intervals. The optically thinner features created in this way were then resolvable by a single measurement.

We could combine these measurements, taking into account the width and diffraction due to the aperture to deduce the total feature created with no aperture present. Rather than accounting for the aperture effects directly however, we use a simpler method. We take the feature shape from the average of the individual measurements. The total absorption cross-section of the feature is obtained by broadening the feature using an electric field gradient, as in preparation for performing a gradient echo. This smears out the absorption to be within our measurement sensitivity.

Specifically we take an average of the profiles in Fig. 5.13a, and scale it to have the same area as the broadened profile in Fig. 5.13c. In this way we estimate the peak absorption as  $150\text{dB}\pm 10\text{dB}$ . The uncertainty in this case is mainly due to the uncertainty in the area of the broadened feature. It is worth repeating that the feature measured here is what remains after selecting for Stark shift as in section 5.4. Without this step, the feature is roughly twice as absorbing (and correspondingly harder to measure accurately).

Some variation in peak height can be seen as a function of position, as is particularly clear when looking at the broadened spectrum. This height variation is not consistent with the light intensity variation across the sample, which is close to uniform and peaks in the middle. It is rather thought to be due to an actual variation in ion spectral density.

Evidence of this spatially varying density is given by similar measurements on a different, 1cm crystal from the same source<sup>2</sup>. These measurements are of the feature height after only a relatively short time ( $\sim 10\text{ms}$ ) with a relatively weak burn-back beam such that less than 10dB of absorption resulted. This was to be far from peak absorption and thus in the linear regime, so that the burn-in rate should then be proportional to light intensity and number density. After dividing the accumulated absorption by  $I(x)$ , where  $I(x)$  is the power measured through the aperture at each crystal position  $x$ , a slow

---

<sup>2</sup>That is, the crystal was also from Scientific Materials, and was likely a piece from the same ingot.

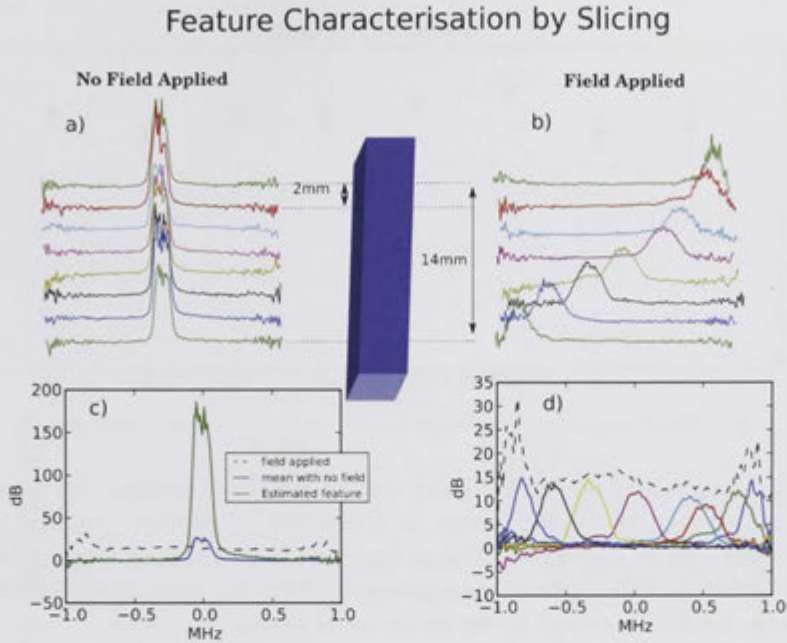


Figure 5.13: Characterising an optically thick feature by slices. a) The burnt-back features as a 1mm aperture is scanned through the burn-back path. b) The same features with an electric field gradient applied. c) A comparison to the case with the aperture removed. d) By using the measured profile of the feature at these positions, and using the integrated absorption measured when a field is applied, the profile of the unbroadened feature is inferred as shown.

variation remains (Fig. 5.14).

## 5.4 Monotonic Stark Shift in PrYSO

To achieve the field gradient required for a gradient echo (based on Stark shift), perhaps the most obvious electrode design is a simple quadrupole as in Fig. 5.15. The resultant Stark-shift is complicated by the crystal symmetry. The measurements show a single anti-hole splitting into up to 4 components depending on the voltage applied to the electrodes. Thus applying a field gradient may make 4 different detuning gradients, which will interfere with the echo efficiency.

The Stark-shift depends on the field direction relative to the transition's static dipole moment, and there are 4 different alignments in PrYSO. This is due to the crystal's  $C_{2h}$  symmetry. The orientations are related by a rotation and a reflection, so that they make an "X". One symmetry axis of the X corresponds to the alignment of the probe beam, which is the  $C_2$  axis of the crystal.

By applying a uniform field as in Fig. 5.15, the different components can be addressed spectrally so that unwanted components may simply be hole-burned away. In this way a




Figure 5.14: Variation of spectral density within a PrYSO crystal. By using a moving aperture, and accounting for the power as a function of position, we can estimate the density of absorbers in the material as a function of position so long as we remain in the linear burn-in regime. This is for a 1cm (along the  $C_2$ -axis) 0.005% PrYSO crystal, different from the 2cm one used in the other measurements.

monotonic gradient can be obtained. The reduction in optical depth however is drastic: to keep only one stark component means a 3/4 reduction in optical depth. We can mitigate this somewhat by retaining two features, if we ensure they stark shift together when the broadening is applied. For this we aim to achieve a field gradient which has components along only a single axis, lying in one of the crystal's symmetry planes ( $x$ - $z$  or  $y$ - $z$  in Fig. 5.15). To avoid impacting the echo, the two features must shift together to a level comparable to the feature sharpness or storage time required, that is the order of 10kHz. Thus the field needs to be very precise.

This symmetry requirement is met for our simple quadrupole shown, however early experiments using these failed to reach the efficiencies expected. Investigating the stark shift caused by these, we find that the monotonicity of the detuning gradient produced is very sensitive to beam alignment and sample position. A new electrode design is made to better suit the crystal. We study the shift due to the new design as a function of position experimentally. By comparison with a finite-element computer model, we can calculate a worst-case for the effects of misalignment.

#### 5.4.1 Spatial Stark Shift from a Simple Quadrupole

The main reason for the poor performance of simple electrode arrangement (Fig. 5.17) can already be seen from the simple measurement in Fig. 5.15: the crystal is much more sensitive to fields along the  $C_2$  axis, while the holder creates a field perpendicular to this. Thus a large field must be used, and only a small misalignment is needed to give a small field but large shift along the  $C_2$  axis.

To better understand the field in the sample, and thus to aid in designing a better set




Figure 5.15: By applying a roughly DC field to the crystal, an initial feature splits into 2 if the field is along along a symmetry axis, or 4 otherwise. Unwanted stark groups are then spectrally resolved and may be hole-burned away. Shown here are 6 electrode configurations (colour coded according to voltage as blue=negative, red=positive, black=ground). This indicates that the dipole moment lies close to the x-axis (the  $C_2$  axis of the crystal), although a possible anisotropy of the permittivity is not accounted for here.

of electrodes, the shift caused by these electrodes was studied as a function of position along the propagation axis. A narrow anti-hole was created within an empty trench using a simple burn-back scheme, with the burn-back beam only applied to a small slice of the crystal using an aperture (as in section 5.3.6). The detuning of the feature when a field is applied is measured by the transmission of a pulse and heterodyne detection. The crystal was deliberately moved to one side by 1mm to be sure of resolving the splitting and better quantify the effect.

Figures 5.16 a) and b) show the dependence of the shift (which is ideally a single gradient), with and without half the subgroups burned away. The maximum splitting is more than 300kHz. Assuming a splitting roughly linear with lateral displacement of the crystal, an alignment accuracy of better than  $30\mu\text{m}$  will be required to reduce the separation below 10kHz. Even without the practical difficulty, this is smaller than the beam diameter. Thus in our present regime (aiming for a storage time of a few  $\mu\text{s}$ ), it is impossible to achieve an efficient echo this way. Using larger bandwidths and shorter storage times would be required for an efficient memory using this holder. Our approach however is to design a better holder.

In the interests of understanding the crystal properties, an attempt was made to fit the above data by comparison to a computer model of the crystal and holder. This could theoretically be used to determine the exact orientation of the static dipoles, as well as the dielectric constant of the crystal. An unique fit was not achieved due to primarily to computational complexity caused by the misalignment and corresponding lack of symmetry, and additional degrees of freedom to account for a slight misalignment of the beam itself (i.e. not going straight through the centre of the crystal).

#### 5.4.2 Stark Shift due to a Rotated Block-Quadrupole

To eliminate the sensitivity to alignment, a different holder was designed. The first consideration was to produce the gradient along the  $C_2$  axis. This reduces both the field strength required to achieve the shift, and the sensitivity to stray, off-axis fields. This is simply achieved by rotating the original electrodes  $45^\circ$ .

Beyond this there is a limit to how much can be done to reduce the sensitivity further, due the requirement that the divergence of the electric field vanish in the absence of charges. This means that the greater the gradient made along the  $C_2$  axis, the greater the gradient along the others must be. We can however minimise the dependence on crystal position and other effects due to dielectrics. The design arrived at here used large square shaped electrodes as shown with holes for optical access.

Further considerations were the dielectric constant of the crystal and the material on which it sits. Symmetric Teflon discs were used to hold the crystal it in place due to its low dielectric constant ( $\epsilon = 2.1$ ). Roughly speaking, increasing the plate electrodes' size compared to the distance between them creates a less linear space dependence, but also decreases sensitivity to misalignment. The design shown (Fig. 5.18) was arrived at






Figure 5.16: Stark-shift as a function of position caused by the simple quadrupole electrodes. A 1cm crystal (and laser) is deliberately displaced by 1mm. Each vertical slice represents an absorption spectrum of an anti-hole created within 0.5mm region. a) and b) show the gradient configuration (Fig. 5.15, LHS) with and without half the Stark-groups burned away. In the lower graphs, the left-hand electrodes are grounded, while the other two have either opposite (c) or the same (d) polarity. In all measurements the voltage applied to the ( shown as red(+)) or blue(-) in Fig. 5.15)  $\pm 35\text{V}$ .

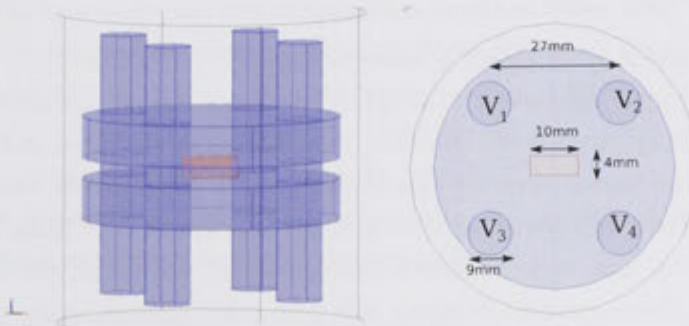


Figure 5.17: The dimensions and details of a simple quadrupole holder, as used in the finite-element model. This design was used in Fig. 5.15 and Fig. 5.16.




Figure 5.18: The dimensions of the rotated quadrupole electrodes and holder, as used in the finite-element model. The pink object is the crystal. The cylindrical objects are Teflon spacers. Surrounding these are 4 copper block electrodes, with holes and slots for optical access. This design was used in for the measurements in Fig. 5.19 and Fig. 5.20.

non-systematically, but appears to be a reasonable compromise between the requirements.

To ensure the correct alignment of the probe beam and reduce distortion of the field, only small 1mm apertures were drilled in the end electrodes to allow the light through. Similarly the crystal position was fixed by precision machining a guide in the Teflon to fit the crystal. With this, the focused probe beam cannot deviate more than 0.5mm from the centre of the crystal.

The measurement of Stark shift as a function of position was repeated on this (2cm) crystal, with the two ‘side’ electrodes ( $V_2, V_4$ ) grounded. A feature was created as before, after which the end electrodes ( $V_1, V_3$ ) were set to  $\pm 25V$ , and the probe beam was scanned from 0.5→2MHz to remove the unwanted components.  $V_1$  and  $V_3$  were then both set to +25V and a further measurement taken. No splitting of the remaining two features could then be resolved. In these measurements we can also see a faint anti-hole probably created by noise on the laser, which has been stark-shifted into our window of interest.

This data was used to match a computer model of this holder, in an attempt to determine the crystal parameters. The static dipole orientation and magnitudes, as well as the (tri-axial) crystal dielectric constant, were allowed to vary, to maximise a goodness-of-fit given by the least-squares frequency difference of the observed lines to the closest predicted. The best fit corresponded to a dipole moment of 0.13 MHz/V/cm, oriented  $12.2^\circ$  from the  $C_2$  axis, with a relative permittivity of  $\approx 10$  along the  $C_2$  axis and  $\approx 4$  perpendicular. These numbers compare well with other measurements in PrYSO using Stark-modulated photon echoes[141], although the dielectric constant was not involved there.

These measurements give us confidence in our model. This predicts a worst-case splitting of 30kHz due to the possible misalignment allowed by the holder geometry, when a total broadening of 2MHz is applied (see Fig. 5.20 and caption).




Figure 5.19: Measurement of the Stark shift vs position with the block-holder, and the 2cm crystal, along with a theoretical fit. One Stark subgroup has been erased by first applying a DC field and burning away one feature (see text). Top:  $V_{2,4} = -25V$ ,  $V_{1,3} = 25V$ . Bottom:  $V_{2,4} = 0$ ,  $V_{1,3} = \pm 25V$ . An additional faint absorption feature is also present. This is probably created by noise on the laser.

An issue with our design is that it was optimised for a 1cm long crystal<sup>3</sup>. Our experiments however used a 2cm crystal, which suffered from non-monotonic stark shift near its front and back faces (as the field goes through a turning point). However, since only a 15mm slot was made for burn-back from the side, the parts of the crystal so affected couldn't be accessed.

## 5.5 Efficient Gradient Echoes

Putting together the considerations above allowed the demonstration of efficient echoes. We use the techniques in section 5.3 to create and analyse a feature, and the electrodes and model used in section 5.4 to broaden it and predict the shift. The layout and stability considerations described in section 5.1 were also applied.

Theoretical modelling of these echoes agrees well, and provides a good basis for discussing limitations and improvements to these experiments.

### 5.5.1 Preparation Sequence

The crystal a 4x4x20mm (D1xD2xb) 0.005% PrYSO single crystal with all faces polished, from Scientific Materials. It was mounted in the block-electrodes as in Fig. 5.18. It was

<sup>3</sup>the very one which was lost in polishing, never to be seen again.




Figure 5.20: The field and resultant Stark-shift predicted by a finite element model of the rotated quadrupoles. The lines show the maximum misalignments allowed by the holder geometry. Left: a simple 0.5mm sideways or vertical translation of the beam. Right: a diagonal misalignment such that the beam is offset +0.5mm at the input and -0.5mm at the output (giving 2.4 degrees). The worst splitting is for the former situation, at 30kHz. The sudden deviations that can be seen at  $\pm 1$ cm are due to the edge of the crystal.

prepared using two beams as in section 5.3. The probe had a peak power of 1.5mW, while the burn-back had a peak power of 3mW, but expanded to a large beam diameter of 2cm.

The spectral trench was created in the centre of the absorption line by scanning both beams at full power over 7MHz for 3 seconds. Next, an approximately 140kHz wide feature was created in the centre of the trench by alternating between burn-back and sharpening steps every second for 10 seconds.

The burn-back consisted of a weak constant illumination from the side while switching the frequency every 10ms between  $\pm 10.2$ MHz,  $\pm 17.3$ MHz, and  $\pm 27.5$ MHz. During the burn-back periods, a 1V p-p modulation was applied to the front and rear electrodes ( $V_1$  and  $V_3$ , opposite voltage on each) to broaden the feature to around 150kHz. This modulation amplitude was tweaked by hand while burning back in only a small section of crystal using an aperture, and observing the transmission spectrum of a small pulse. The sharpening step used the side-beam, alternating between a scan from -175kHz to -75kHz, and 75kHz to 175kHz. The scan rate was  $1\text{MHzs}^{-1}$ .

To select for Stark-shift direction, an approximately uniform electric field was applied with -20V and +20V to the front and back electrode respectively ( $V_{3,1} = \pm 20\text{V}$ ) whilst grounding the side electrodes. The probe beam was scanned from 200kHz to 1.4MHz to burn away one of resulting features.

The electrodes were again grounded, and a ‘clean-up’ step was applied to the low frequency side of the feature. The probe beam was scanned weakly between -2MHz and -150kHz at  $2\text{MHz/s}$  for 1s.

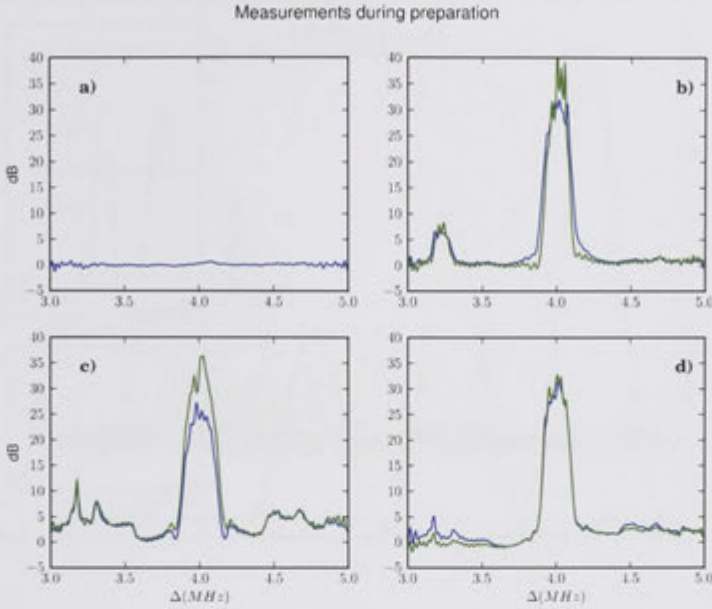


Figure 5.21: Absorption measurements during the preparation sequence leading to efficient echoes. a) The spectral trench, 10ms after it was created. b) A feature is created using sequential burn-back (blue line) then clean-up near the edge of the feature (green line). The peak optical depth is already beyond the sensitivity of the measurement. c) The feature just before selecting for Stark components. A relatively large background has accumulated mostly due to noise on the laser. d) A DC field is applied and one Stark group is burned away, reducing the total absorption by approximately half. The actual burning step was not recorded. A clean-up field was then applied to the low-frequency side of the feature and recorded in two steps (blue then green). See text for details.

Periodically throughout the preparation sequence, a quick absorption measurement was made using a short, weak pulse along the probe path and heterodyne detection. These measurements are given in Fig. 5.21.

### 5.5.2 Efficiency Measurements

The finished anti-hole was broadened by applying +20V to the front and back electrodes, while the side electrodes remained grounded. A  $0.6\mu\text{s}$  (FWHM, intensity) Gaussian pulse was then input along the probe path. The intensity of the probe beam was  $\sim 1\mu\text{W}$ , chosen to be much less than a  $\pi$  pulse but bright enough to give a good signal. After a delay, the applied voltage was then switched to -20V. These storage and retrieval steps were repeated 5 times with a delay of 10ms between repetitions. The retrieved echoes were measured using balanced heterodyne detection.

The entire sequence was repeated several times before data was recorded to avoid burn-in effects and ensure repeatability.




Figure 5.22: Efficient gradient-echoes. The main figure shows input and retrieved intensities recovered from the heterodyne measurements, each consisting of 5 averaged traces. The simulated echoes use parameters determined from the inset panels. These show the simulated and actual measurements of the initial (upper) and broadened (lower) feature. The simulations mimic the limited spectral resolution of the actual measurements. The additional delay of the measured echo compared to the simulation may be due to off-resonant interaction with the edges of the spectral trench, which were not modelled. The interaction with these edges is larger during the echo than the input as the field gradient shifts the edges closer. The unusual absorption profile of the broadened sample is thought to be due to spatial variation in the crystal.

The resultant echoes (Fig. 5.22) were measured using balanced heterodyne, and compared to ‘input’ pulses which were applied 10ms after the initial trench-burning step. These serve as references to account for optical and detection losses. The efficiency was measured as a ratio of the area (energy) of the echo to that of the input. For the closest echo, there is a small time overlap between the input and retrieved pulses. In this case, the retrieved area is taken to start at the point in time where 99% of the input pulse had already passed. The efficiencies determined are  $69 \pm 2\%$  for the first and  $45 \pm 2\%$  for the longer-delayed echo. The uncertainty accounts for the uncertainty in the transmission of the input pulse through the trench.

To model the experiment, a spatially uniform feature shown in Fig. 5.22b) was used. We assume this to exist over the 14mm of crystal depth illuminated by the burn-back beam. A finite-element model of the electrodes (Fig. 5.20) was used to predict the broadening. The major factors not accounted for in the model are the spatial variation of the initial feature (due both to a non-uniform crystal and a non-uniform preparation beam), and other spectral structure such as the edges of the trench.




Figure 5.23: The experiment was also performed using direct intensity detection for several delays. The largest efficiency seen here is  $71 \pm 5\%$ , with most uncertainty coming from measurement of the input pulse size. As the echoes here are not measurably larger than for the heterodyne case, it is indicated that any distortion of the spatial mode due to storage must be similarly small. Small noise spikes can be seen caused by electrical pickup from the oscilloscope trigger. These occur 100ns before the field is switched.

The experiment was also repeated using similar preparation and plain intensity detection (Fig. 5.23), as any distortions of the beam mode due to storage may ruin heterodyne mode matching. The sensitivity in this case is much lower, but the measured efficiencies are not significantly different. We thus expect there is little distortion of the spatial mode due to the storage.

### 5.5.3 Predicted Improvements

We use our model to predict the effect of changing experiment parameters. By simply scaling the experiment so that we have 40mm of active crystal depth (instead of 14mm), an efficiency of 88% is predicted. The limitation here is the estimated background absorption of 2% in the current experiment, scaled to 7%. Further scaling soon begins to decrease the efficiency as this background becomes more significant. As described in section 5.3.1 however, significantly better should be possible. In the limit of perfect hole-burning, the modelled experiment reaches a peak efficiency of around 93% when scaled to 10cm of absorbing crystal length. Further improvements would require a sharper-edged feature. These two simulated situations are given in Fig. 5.24. The increased delay and temporal width is due to off-resonant interaction with the high-detuning side of the spectrum as shown in the inset.




Figure 5.24: Predicted echoes by simply scaling the experiment to 4cm (dashed), and by scaling the experiment to 10cm while also assuming no background absorption (dotted). This figure follows from Fig. 5.22

## 5.6 Quantum noise measurements

We have so far demonstrated an efficient classical memory. To consider practicality for quantum storage, as discussed in chapter 2, we need to make measurements sensitive at the quantum level. In particular we make repeated homodyne measurements of (nearly) coherent pulses after storage in the memory.

Each homodyne measurement corresponds to the quadrature operator, and we expect to see variation as the coherent states are not eigenstates of this operator. The distribution of measured values maps out the quantum state.

Efforts at similar measurements usually make 10s of thousands of measurements over all phases, but otherwise identical conditions. This allows the detailed reconstruction of the quantum state using tomographic, maximum-likelihood, or otherwise based techniques[160, 35]. If we are only concerned with Gaussian states, we only need to make (repeated) measurements at 3 different phases to determine the covariance matrix.

Due to limitations in our detection system, sets of only 1000 measurements could be made at one time before re-preparing the feature. Slight variations in the preparation meant we are measuring a slightly different memory each time. A further difficulty was our unstable interferometer meant measurements could only be made at random phases. Thus while we are interested only in Gaussian states, we cannot measure a single phase at a time to implement the normal procedure. Instead we use a maximum likelihood technique to fit a Gaussian state to a set of 1000 measurements. We combine many sets



to improve the precision. The parameters of these states can then be compared to the benchmarks derived in chapter 2.

### 5.6.1 Quadrature Measurements

To obtain a ‘quorum’<sup>4</sup> of quadrature measurements, an echo experiment was performed similarly as for efficiency measurements. The local oscillator was set to the same frequency as the probe beam (for homodyne), and the measurements are of amplitude rather than intensity. The noise minimisation considerations discussed in section 5.1 were crucial to these measurements.

Measurements at the quantum level are more easily performed on weak signals, where classical noise is insignificant. Such weak input pulses were made using large passive attenuation ( $\sim 30\text{dB}$ ) in the RF driving the double-pass AOM in the probe path. Small changes could then be made to the RF power digitally. After the feature was made, 1000 such pulses were input and recalled with a delay of 1 ms between each.

For each pulse, a  $40\mu\text{s}$  time window was recorded for post analysis. A single such measurement with the largest pulse size used is shown in Fig. 5.25. This is at the maximum phase for the retrieved pulse. The different time regions of this trace used in post-processing are highlighted.

The phase reference pulses (section 5.1.4) were large Gaussian profiles to keep their frequency spectra narrow, with FWHM of  $10\mu\text{s}$ . They were detuned from the input pulse by 2.4MHz, to be outside the memory bandwidth, but still within the transmissive trench window. The actual pulses to be stored/measured were timed to be far away from the initial reference, avoiding any FID caused by it.

Small time windows with no input before and after the pulse to be measured were used as references for the DC level on which the input pulse appears. This is an effective high-pass filter, filtering as much low frequency noise as possible while impacting the signal as little as possible.

The temporal mode envelope for the input and retrieved modes was measured by finding the peak phase, and averaging together many measurements taken near this phase value (Fig. 5.25, inset). To obtain a quadrature sample, the homodyne trace is multiplied by this mode function, then the area is taken. The average size and difference between the reference pulse amplitude and phase was used to filter out samples where the amplitude or phase was changing rapidly ( $< 5\%$  of samples). In principle, the amplitude of these pulses could further account for laser amplitude fluctuations by scaling the measured quadrature value. Unfortunately, burn-in effects from these references (section 5.6.5) made such a scheme unstable.

The experiment was repeated with 4 different pulse sizes, with the smallest chosen to be near the detection noise level, and the largest chosen to be easily resolvable.

---

<sup>4</sup>A set of measurements spanning the set.




Figure 5.25: The parts of a homodyne trace used for measurements of quadrature noise introduced by the memory (with a relatively large input pulse for clarity). The main graph shows many averaged measurements within a 5 degrees of the LO phase at which the homodyne signal is maximised. i),ii) The input and retrieved pulses. iii) The scaled, smoothed mode envelopes used to make quadrature measurements. iv) Sections used as a DC reference to cancel slow local oscillator variations, effectively a time-domain filter. v) Large reference pulses before and after were fit to determine the phase of the measurement (section 5.1.4).

Saving this data is a time consuming process, such that after each set of 1000 samples, the absorption feature had to be recreated. 5-10 sets of 1000 samples were recorded for each pulse size. Immediately after these sets were recorded for a given size, the experiment was repeated an additional 5-10 times, but with the trench burning step continued in place of the feature preparation steps. This was to give a measure of the input pulse under the same conditions for electronic noise and reflection losses. This entire process was significantly complicated by the time needed to record this data. Sets were only kept if the laser remained locked and the helium level and flow rate remained constant.

### 5.6.2 Maximum Likelihood Estimation of Covariance

Given a set of 1000 quadrature measurements at random phases, we aim to find what kind of state produced them. We assume that additional noises due to the local oscillator and detection electronics are independent and closely Gaussian. We expect the same for any noise that is added by the memory. The signal observed at the detectors should then be qualitatively the same as that from a perfectly detected, but noisy Gaussian state. We aim here to find the parameters of this state, then work backwards to estimate the effect of the memory.

The measured quadratures should map out a sinusoid in the mean, with a variance that also varies sinusoidally [29].

$$\langle X_\theta \rangle = X_0 \sin(\theta - \theta_0) \quad (5.1)$$

$$V_\theta = V_1 \sin^2(\theta - \theta_0) + V_2 \cos^2(\theta - \theta_0) - V_{12} \sin^2(\theta - \theta_0) \quad (5.2)$$

5 parameters can thus completely determine a Gaussian state- two are the average amplitude and phase,  $X_\theta$  and  $\theta_0$ , while the remaining three describe the covariance,  $V_1$ ,  $V_2$ ,  $V_{12}$ .

The likelihood of a given state producing a quadrature value at a certain phase is given by the normal distribution centred at the mean  $\langle X_\theta \rangle$  with standard deviation  $\sqrt{V_\theta}$ . To determine the likelihood of a set of measurements, we simply multiply the likelihood of all the individual measurements (in practice we add the logarithms to avoid numerical issues of rounding very small numbers).

To find the state that maximises this likelihood, standard optimisation methods such as simulated annealing may be used, but are sensitive to outliers and potentially to local maxima. Here we first use simulated annealing just to give an initial guess at the parameters, then use an adaptive Metropolis[161] algorithm to fully sample the likelihood distribution of the parameters. This algorithm is closely related to simulated annealing in that parameter space is moved around in random jumps, which are kept based on the likelihood of the new position and a kind of ‘temperature’ giving the probability of keeping a lower likelihood step. The Metropolis algorithm keeps the temperature constant to

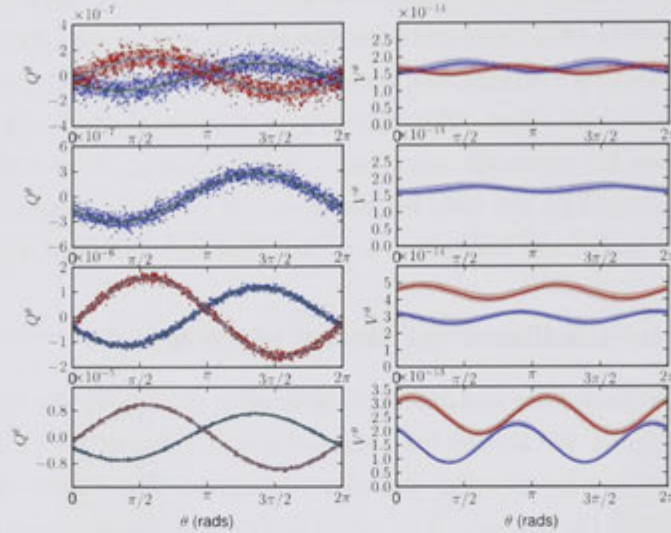


Figure 5.26: a)-d) The first set of 1000 quadrature measurements for each attenuation level. The red(input standard deviation) and green(output standard deviation) shading together with the cyan(input mean) and green(output mean) lines correspond to the fits obtained using the Metropolis-Hastings method. e)-h) Show the determined variance vs phase for the fit, with the width of the shading in this case indicating uncertainty in the variance.

sample the entire space, while simulated annealing lowers it to arrive at the peak. The adaptive version that we use optimises this temperature as it runs, and requires very little user intervention<sup>5</sup>.

The fitted value of each parameter is given by the mean of its distribution. This is not necessarily the most likely value (although we could also find that), but rather takes into account peakedness in the distribution and avoids local maxima. We can also use the measured parameter distribution to obtain the full covariance matrix for parameters.

To fit each of our measured data sets, the adaptive algorithm was run for 10000 accepted steps as ‘burn-in’, i.e. to let the algorithm adapt. A further 50000 steps were taken, with only every 5th sample kept to give 10000 samples, each sample consisting of a value for all 5 parameters. The uncertainty in each parameter is taken as the standard deviation of its sampled values.

Figure 5.26 shows the first set obtained for each pulse size.

<sup>5</sup>Perhaps the biggest advantage of removing tunable parameters is to remove the temptation of tuning them until results that look “good” to the experimenter are found.

### 5.6.3 Attenuation and Additional Noise

The state we measure at our detectors is attenuated by optical and detector losses (and after storage by the memory, the imperfect memory efficiency), and is contaminated with electronic noise due to the detectors as well as local oscillator noise. We aim to work backwards to recover the parameters of the state input into and recovered from the memory. We can then determine if the effect of the memory is just an attenuation, or whether it also adds spurious noise.

As we are only interested in variances, and our noise sources are independent, the analysis is remarkably simple. Variances from independent sources add, and attenuation just scales the classical variance linearly. Thus to work backwards, we simply divide by the attenuation and subtract the noise sources.

Specifically, our measurements at a given phase have variance according to

$$V_{\text{meas}} = \eta_{\text{opt}}\eta_{\text{det}}(V_{\text{class}}) + V_{\text{shot}} + V_{\text{elec}} + V_{\text{other}}$$

where the subscripts denote the source of each noise term. The first term denotes the actual value we are measuring, which we may think of as the classical noise on the light at the input port to the detector<sup>6</sup>. The remaining terms represent the quantum noise of a coherent state, electronic noise, and any other source such as uncancelled noise on the local oscillator.

The coefficient  $\eta_{\text{opt}}\eta_{\text{det}}$  indicates linear attenuation of the classical noise, with factors for the optical losses and detection inefficiency.

The amplitude measured is

$$\langle X \rangle_{\text{meas}} = \sqrt{\eta_{\text{eff}}\eta_{\text{opt}}\eta_{\text{det}}} \langle X_{\text{input}} \rangle$$

where  $\eta_{\text{eff}}$  is the memory efficiency.

We wish to determine any noise added during storage. Thus we break down the measured classical variance as

$$V_{\text{class}} = \eta_{\text{eff}}V_{\text{input}} + V_{\text{mem}}.$$

$V_{\text{input}}$  is the classical noise on the input.  $V_{\text{mem}}$  is the value we wish to determine, indicating any additional noise arising from the memory. Measurement of the ‘input’, i.e. when the memory isn’t created, gives us  $V_{\text{input}} = V_{\text{class}}$ .

To determine  $V_{\text{mem}}$ , we must first determine the other parameters. The efficiencies can be obtained easily by classical measurements. To determine the other noise parameters, we consider measurements in the time domain. More typically, this is done in the frequency domain, but we are already performing time-based filtering by subtracting the DC level

<sup>6</sup>More generally it’s the difference from the shot noise. If somehow our detected state was squeezed, this could be negative.

seen before and after the pulses. The consequences of the filtering are clearer in the time domain.

The noise levels were determined by the following measurements.

We perform the experiment as usual, but with the input port of the homodyne detector blocked, giving us a set of time traces similar to Fig. 5.25 but with only noise present. Each is adjusted according to the measured DC level immediately before and after the input and echoes, as usual. Each is then convolved with the normalised envelope of the input mode. The variance of 3000 such traces is taken, and plotted in Fig. 5.27a. Every point in the plot gives the detection noise that would be observed if the input pulse were to arrive at that time. This was repeated for the retrieved mode in Fig. 5.27b. These panels give us  $V_{\text{shot}} + V_{\text{elec}} + V_{\text{other}}$  for each mode, as a function of time. The measurements were repeated with all light blocked to isolate the contribution of electronic noise. Repeating the analysis gives us  $V_{\text{elec}}$  in panels c) and d).

To isolate the contribution of low frequency detector noise, all the above measurements were re-analysed with the modes artificially shifted in frequency. The envelopes are then complex, modulated in frequency by  $e^{i\omega t}$ . It is done this way so that they remain normalised, and thus include the same shot-noise contribution as the original modes (section 1.2.8).

In Fig. 5.28 we subtract the bottom panels of Fig. 5.27 (electronic noise) from the top. This isolates the noise due to the light alone. At higher frequencies, where the detectors are shot-noise limited, we have a measure of the shot-noise level,  $V_{\text{shot}}$ .

It should be noted that sampling for only a short time, as here, to estimate the DC level is itself inherently limited by shot noise. The level added by this can be estimated as the ratio of the lengths of the pulse mode to that used for estimating the DC level. For the widths used in the displayed measurements, this is  $\approx 0.1V_{\text{shot}}$ . This factor could be reduced by increasing the length of the reference regions, at the expense of filtering less noise. The majority of the extra noise observed at DC is due to low-frequency electronic noise of the detectors.

Thus the shot noise level is measured to be  $7 \times 10^{-7} \pm 0.3V^2\text{s}$ . The base noise level for the input mode is  $1.8 \times 10^{-6} \pm 0.1V^2\text{s}$ , while that for the retrieved mode is  $1.9 \times 10^{-6} \pm 0.1V^2\text{s}$ . The uncertainties in these levels corresponds to a relative error of  $\sqrt{2/N}$  where  $N = 3000$  is the number of shots combined to get the variance.

#### 5.6.4 Comparison with Benchmarks

We can now calculate the added noise for each data set, for comparison with benchmarks described in section 2.7. For the two smaller pulse sizes this is just the variance above the base detection level as any classical noise is negligible. For the two larger sizes we find the variance above the noise expected based on attenuation of the measured input. We then combine the values for each set by taking the weighted mean according to the uncertainties calculated from our fitting method (that is, the contribution of each set is

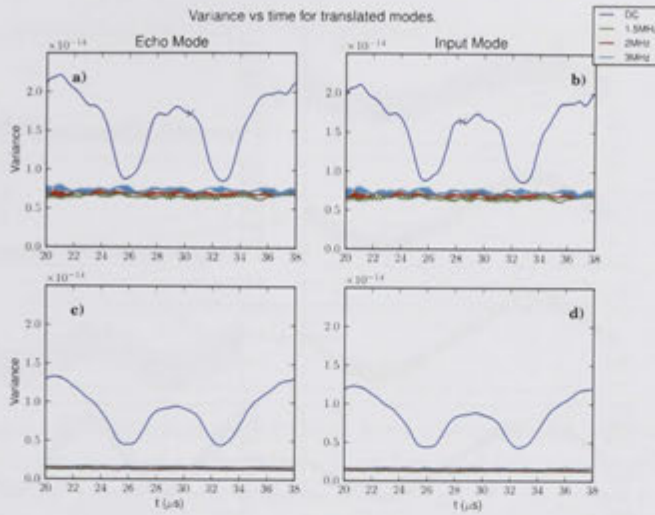


Figure 5.27: The experiment is performed with no input beam, and the noise in a temporal mode defined by the pulse is measured as it is translated in time. This shows the effect of the time-domain filtering, where we subtract off a DC level using a patch before the input, and after the retrieved pulses (where the depressions are). The left side corresponds to the echo mode, and the right to the input mode. The bottom panels are obtained by repeating this experiment with no light incident. The experiment is also repeated again with the modes shifted in frequency, illustrating that much of the variance is due to low frequency noise. Crosses correspond to the times of the actual input/retrieved modes.

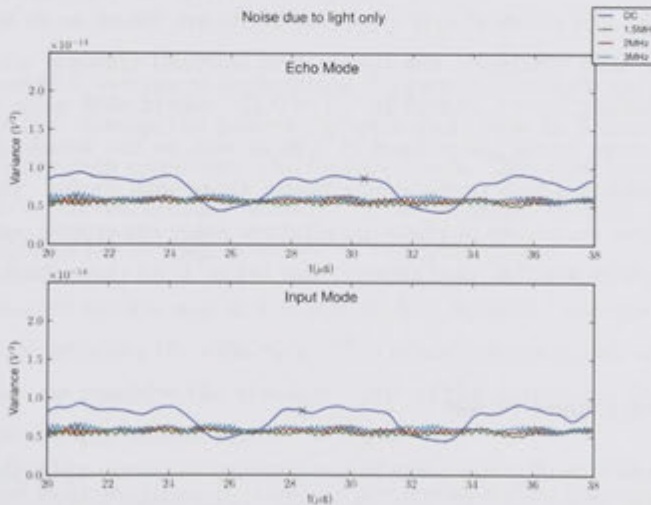


Figure 5.28: Subtracting the variance when no light is incident from that with just the probe blocked reveals the noise on the light. At high frequencies this is the shot-noise level. The dips at DC are caused by the background subtraction intervals. Crosses correspond to the times of the actual input/retrieved modes.

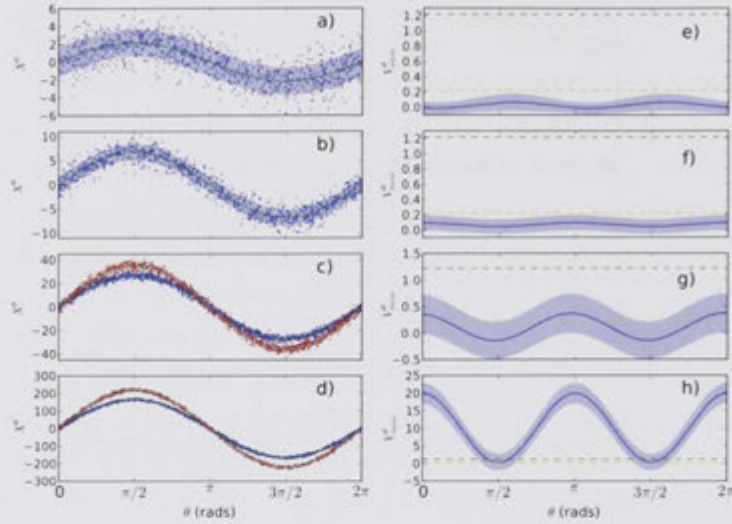


Figure 5.29: Calculated added quadrature noise ( $V_{\text{mem}}$ ) and comparison with benchmarks, normalised to the shot-noise level, and with the memory phase-shift accounted for. a)-d) Example data sets of 1000 and fits to them, for  $\langle N \rangle = 4, 30, 500, 19000$ . e)-h) Calculated added quadrature noise vs phase. The wide blue line indicates the  $1\sigma$  uncertainty. The dotted lines give the phase-insensitive classical (higher) and no-cloning (lower) limits for the memory.

weighted according to the variance of the estimated variance!)

The levels for the classical and quantum limits are found as in section 2.7.2. For a memory of  $\eta = 0.61$  efficiency, the (phase-symmetrical) classical and no-cloning limits correspond to adding  $2\eta = 1.22$  and  $2\eta - 1 = 0.22$  units of shot noise respectively.

The added noise levels are plotted in 5.29 as well as the benchmark levels. For the two smaller inputs, we find performance better than both limits. Measurements of the larger pulses show an excess of phase quadrature noise apparently added by the device, while it is uncertain whether performance was better than these limits for the amplitude quadrature.

### 5.6.5 Added Phase Noise

The apparent addition of phase noise to the memory requires explanation. It is consistent with a problem in our measurement of phase. During the feature preparation process, a small background population is replaced into the trench (by laser noise, or just decay with time). The pulses used as phase references are then transmitted through this slightly absorbing background, in which they burn a hole over time. This spectral hole changes the transmitted pulse shape, and delays them slightly. When measuring the input pulses, the preparation step was replaced by continuous trench-burning. Thus there is no opportunity



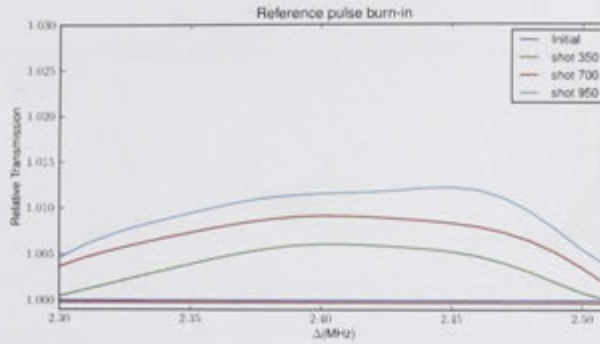


Figure 5.30: Due to a small residual background level, the reference pulses slowly burn in, and become less reliable as phase references. This figure shows the relative transmission spectra of the reference pulses. Each trace is an average of 100 pulses in a single set of 1000, centred at 250, 500, and 750. The transmission is relative to the first 100 pulses of the set.

for a background to accumulate.

The increasing transmission with time is shown in Fig. 5.30.

### 5.6.6 Phase Modulation of the Temporal Mode

As discussed in chapter 4, storage at high optical depth can introduce some distortion into the temporal mode (i.e. change the pulse amplitude and phase as a function of time). A slight phase shift can be seen across our echo for this reason. This shift has the consequence that the echo mode is not represented perfectly by the envelope we used (Fig.5.25), which was calculated simply as the average envelope at the peak phase.

The mode estimated in this way is actually slightly smaller than the real mode, and thus we slightly underestimate the efficiency. This effect is quite small, as can be seen in the Fig. 5.31, where we consider the absolute value of the envelope. The difference will only alter the measured quadrature values by  $\sim 1\%$ .

If this shift were more extreme, it could be accounted for by treating each measurement as giving two orthogonal quadrature samples, but each with appropriately reduced accuracy. The limiting case for a large shift is then of course a heterodyne measurement.

Using our power of hindsight, we note that performing all these measurements deliberately as heterodyne (by detuning the local oscillator) would likely have been much more sensible, due to the large effort required to account for low frequency noise in this setup.




Figure 5.31: The magnitude a) and phase b) of the echo temporal mode. It can be seen that the phase varies across the retrieved pulse. It can be seen however that the vast majority remains in the mode given by the mean envelope.

## 5.7 Other Demonstrations

### 5.7.1 Multiple Stark Groups

In the standard protocol in the optically-thick regime, with only one detuning group, switching the field again results in a small echo (going to zero as the optically thickness is increased). We predicted (section 4.4.5) that if two groups corresponding to opposite detuning gradients are present, we get a small first echo, but a relatively large second one. The potential to manipulate different sub-sets of absorbers gives an extra handle on manipulated light.

To test our predictions experimentally, a feature was prepared as already described, with one set of Stark subgroups removed as normal. An echo was performed as normal with the field switched at  $t = 2\mu\text{s}$ , and then reversed again at  $t = 4\mu\text{s}$ . The results recorded using heterodyne detection. The experiment was then repeated identically, but with the step for burning away one Stark subgroup removed. The results are given in Fig. 5.32.

As expected, leaving all groups intact gives a smaller first echo, but a larger second echo. The significant distortion to the mode however is also comparable to Fig. 4.14.




Figure 5.32: Double switching leads to two stage memory. When one stark feature is burned away, the sample emits most of the light it had absorbed. On switching the field again, little light is released. When two Stark groups are present, the first echo is suppressed. On switching again however, a relatively large echo is released. This was predicted theoretically in section 4.4.5.

### 5.7.2 The Edge Effect

We proposed in the last chapter that in the optically thick limit, the anti-hole effectively hides one side of the spectrum. This was born out by the high efficiency results, where only one side of the feature had its background removed.

As a more explicit example, we perform an echo experiment 3 times.

An anti-hole is made using a long, 30s burn-in sequence, with the background on both sides ‘cleaned’ up. An echo experiment is performed with a positive detuning gradient (detuning increases with propagation depth). Next the sequence is repeated, but the power of the clean-up beam on the high frequency side is set to zero. This allows background absorption to accumulate. The echo is performed as before, but with the field polarity reversed. The experiment is repeated one more time with the same feature shape, but with the polarity set back to normal.

According to our analysis last chapter, we expect the first two experiments to have a similar efficiency, while the third should be reduced. This is seen in Fig. 5.33.

In the first it should not matter which side of the feature is ‘seen’ by the light as both sides are roughly the same. In the second we switched the gradient so that the high frequency side was seen by the light first. The echo efficiency is hardly affected, as we expect, while the transmission is reduced. Finally in the 3rd experiment the light is again incident on the low-frequency side. This time the efficiency is reduced.

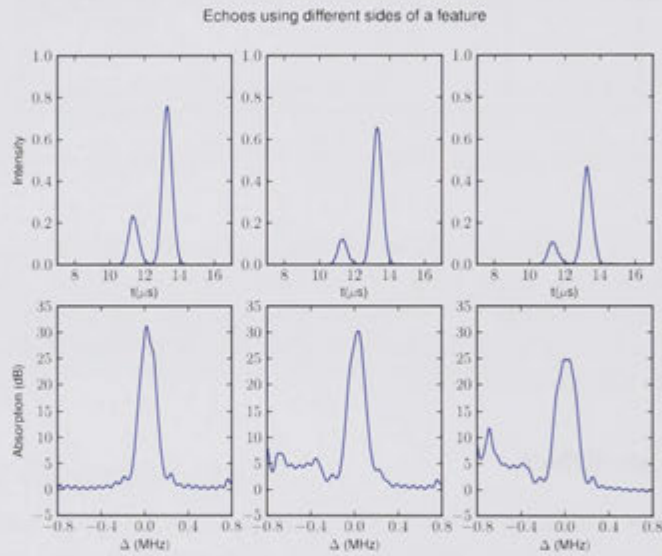


Figure 5.33: The field gradient is applied so that the input side of the crystal is detuned negatively (the gradient is positive with depth). The upper panels give the measured echo, while the bottom panels show a single shot (dynamic range limited) measurement of the initial feature. Left: Both sides of the feature are 'cleaned'. Middle: The high detuning side of the feature is cleaned, and the field polarity is reversed. Right: The high detuning side is cleaned, and the field polarity is returned.



---

# Prospects for Practical Quantum Memory using Rare Earths

---

In the previous chapters we have discussed the requirements for implementing high performance, practical quantum memory. We focused on REIs in the solid state for potential to maximise the main requirements of strong coupling, long storage time and high density. Directions to realise this potential are the focus of this chapter.

The experiment in the previous chapter represented the first demonstration of a quantum memory in a solid, and the first in any system to perform unconditionally better than the no-cloning limit. It was not, however, a demonstration of a practical device. Still higher efficiencies and the storage of many more modes will be required. Further work is also required to develop supporting devices, but we focus on memory.

Significant improvements may be realised by extending the experiment, including greatly extending the storage time, before running into bandwidth or background absorption limits due to our use of hole-burning. By using different materials, the ceiling may be pushed back, but remains limited. The fact remains that our technique wastes most of the available interaction strength, and thus the TBP, that is possible in a REI doped crystal.

If we can make better use of the available absorption, the theoretically achievable efficiencies and mode-capacities compare very favourably to the requirements of quantum repeaters, and even quantum computers. We offer three approaches to achieve this in current materials. The approaches are multiplexing many memories similar to our demonstration, using optical rephasing, and using a Raman absorption with a spin polarised sample. Taking a step further, we consider a different class of low-disorder materials that avoid doping in favour of stoichiometry. This is a particularly exciting direction not only to greatly increase mode capacities, but to simplify experiments, and even approach the physical limits of light-ensemble interaction strength.

Once memories performing at a practical level are demonstrated, probably the largest obstacle to their use will be interoperability with other devices, in particular quantum light sources. We highlight current efforts. As quantum light sources use similar techniques to quantum memory, it is likely that these will be available and improve in conjunction with

memory.

## 6.1 Improvements to the Current Experiment

We discussed the limitations of hole-burning in chapter 3. In particular, the materials that hole-burn well are limited by the hyperfine splittings. We may push these limitations a little further yet, and there may be benefit to doing so. Using other materials, particularly those based on Europium may gain us another factor of  $\sim 100$  in TBP.

Even if performance falls short of practical requirements, these materials still present a strong case for testing the long term storage of quantum states. The world record for storing light is still held by Longdell et al. using a feature burned in PrYSO as here[1], and indeed a gradient echo with (potentially long-term) shelving has also been performed[69]. Highly efficient, long-term quantum storage is also of interest from a fundamental point of view.

### 6.1.1 Improving Efficiency and Time-Bandwidth Product

With a coherence time of  $\mathcal{T} \approx 100\mu\text{s}$  and a maximum bandwidth of  $\mathcal{B} \approx 10\text{MHz}$ , the theoretically achievable TBP in PrYSO is limited to  $(\mathcal{T}\mathcal{B})_{max} \approx 1000$ . Considering the storage of Gaussian pulses 95% within the memory bandwidth, this translates to roughly 10 ( $2\sigma$ ) pulse widths before the total efficiency is reduced to 90%, or 5 well separated pulses. These numbers don't account for off-resonant absorption from ions outside our trench, which will probably limit it further. This number of modes is significantly less than required for practical use.

Of the materials we surveyed in chapter 3, those based using the  ${}^7F_0 \rightarrow {}^5D_0$  transition of  $\text{Eu}^{3+}$  offered the most potential. In EuYSO, a coherence time of 2ms combined with a maximum bandwidth of around 50MHz gives a TBP of  $10^5$ , around 1000 pulse-widths above 90%, or 500 well separated pulses. Again we haven't considered off-resonant absorption, which will be a factor when the optical depth is high. The major difficulty with Eu based materials is the much lower oscillator strength- less than 1/10th of  $\text{Pr}^{3+}$  in YSO. To achieve the required OD, an  $\text{Eu}^{3+}$  material will undoubtedly need a cavity or multi-pass cell. This system may well be enough for some repeater applications if combined with long term shelving.

Of course it may be possible to find more suited materials (likely with  $\text{Eu}^{3+}$  as a dopant) with larger nuclear-quadrupole moments and thus higher potential bandwidths, without affecting the hole-burning and coherence times. It is not yet clear how far this can go.

### 6.1.2 Long-term Shelving

The potential to store quantum information on the spin states for very long periods is one of the main draw cards of the rare earths, and solids generally. PrYSO still has the longest

reported light storage time (2s)[1] demonstrated in any system. A sample of PrYSO was also used to obtain much longer coherence times (30s)[122] for certain inputs[140]. Yet there is room for much improvement using materials such as EuYSO.

To add spin-shelving to our current experiment, the first difficulty will be to create an empty shelving level, and to only use a single transition for the main absorption peak. If the transition is chosen well, this may only reduce the optical depth by a factor of 3. Simply transferring the coherence to a spin state (and rephasing its inhomogeneous broadening) should achieve storage times according to the ground state  $T_2$ ,  $500\mu\text{s}$ [162].

To extend the time to 1s, a precise magnetic field is needed to achieve the conditions for ZFOZ[163, 138]. This means splitting the level and the available optical depth yet again. The hole-burning should get easier though, as the field may reduce cross-relaxations. To extend this further, DDC techniques and very precise RF pulses will be needed [163, 122].

Whether we can easily achieve the required optical depths in this configuration is unclear. If hole-burning is eased, there is a good chance. A lot will also depend on which critical point is chosen for the ZFOZ technique as this will determine the splittings and the bandwidth available.

If an efficient memory can be realised in EuYSO, shelving times should be significantly better due to its magnetic moment-less ( $J = 0$ ) ground state. The spin Hamiltonian has also been characterised, and the turning points are much flatter[134], so that ZFOZ should work better. The main difficulty (beyond the achieving the optical depth required for efficiency) is the strength of the fields required, of a few Tesla. For fundamental tests of quantum coherence, EuYSO is easily the best choice of material. It is not out of the question to obtain coherence times of minutes or even hours<sup>1</sup>.

## 6.2 Techniques for a Real Device

We can as accurately predict the required parameters of future quantum information devices as we can predict the future uses of it— which is to say not very accurately. Yet we can safely say that quantum memories with higher TBPs, overall efficiencies, and storage times will be useful for both repeaters and quantum computers.

We discussed in chapter 2 the requirements of a quantum repeater. We concluded that we need storage times at a bare minimum of 10ms (time for a few signals transmitted over 1000km). At this scale, we need to have efficient, shuffle-able modes at a number comparable to the loss rate of the entire link, including transmission, detection, and other losses. A further requirement is efficient quantum light sources to match. We may hazard a guess at 1000 modes required, expecting to be out by an order of magnitude or so. As the storage times go up, these requirements relax accordingly, although the number of modes remains proportional to the transmission rate.

---

<sup>1</sup>Long enough for instance to test relativistic effects by driving one half of a stored entangled pair around the block!



The major, but more distant application of quantum information is quantum computing. Precise requirements here are again difficult to guess, due partly due to the number of different physical realisations, and partly due to the number of different scales of problem. It's clear that the memory must have a long storage time compared to that of a gate operation, and be able to keep up with the operation rate, that is a large bandwidth. Overall, a large TBP. Efficiencies must also be high, the requirement increasing with the number of operations to be performed. Error correcting protocols may ease this, but can't even begin operating with less than 50% efficiency[164]. We may also hazard to guess that for true scalable quantum computers, the physical size of a memory device will be important to allow on-chip type designs. If operations are fast enough, it's possible that long-term shelving of quantum information —the “hard-disk” of classical computers— will not be necessary in this context.

To create an optimum memory for quantum computing applications then, we must make full use of the available optical depth in a REI material. The density of absorbers in such materials is one of their biggest draw cards, but only if we use it. Ideally we would increase the density of absorbers further using a stoichiometric, not doped crystal. It seems feasible to grow such crystals with inhomogeneous broadening lower than the hyperfine splittings, allowing huge optical depths, a much simplified experimental procedure and no fundamental limit on achievable bandwidth. It should also be possible to make use of the entire inhomogeneous line of the long-studied doped materials, by using a technique based on optical rephasing.

For a repeater we require long storage times, but the longest coherence-time of an optical transition reported in a solid is that of Erbium, at  $\approx 6$ ms. Thus a useful repeater will likely require spin-shelving. We offer two possibilities to create a high mode capacity memory with long-lived shelving. The first is the obvious, but surprisingly powerful one of creating more mode-capacity simply by using many memories. We can do this either by preparing many spectral regions within the inhomogeneous line, or by multiplexing spatially. The second possibility is that if we can prepare an entire inhomogeneous line in the same hyperfine state, it should be possible to use a virtual transition combined with a gradient echo technique to remove the bandwidth limitations of hole-burning. The ideal repeater material however remains as for computing: if we can achieve hyperfine-resolved, stoichiometric materials.

### **6.2.1 Highly Parallel Hole-burning Memories**

Given present and past experiments, it appears that our hole-burning strategy can achieve high efficiencies and long storage times at the level required for quantum repeaters, just only for a few temporal modes. A simple, and potentially very powerful way to boost the mode capacity is to create many parallel memories, using the same medium.

In frequency space, one can create many spectral windows within the inhomogeneous line, each containing a memory at a different frequency. In PrYSO the inhomogeneous

line is  $\sim 10\text{GHz}$ . Given the total span of its hyperfine structure, we may imagine making a memory every  $30\text{MHz}$ , giving on the order of 100 memories. These memories may be controlled simultaneously with the same external field (for CRIB), or alternatively by using separate beams to address each memory independently (for any technique with shelving). Probably the biggest limitations here are the resources required to prepare and maintain such a structure, as well as the optical control to rapidly perform the shelving step for all these memories.

More powerful is the possibility to use many spatial modes. Ensemble techniques essentially store a 2D image which might be decomposed into spatial modes, e.g. in the form of pixels. Using all these memory ‘pixels’ in parallel may be quite straightforward—in the first instance all we need is to replace our single detector by a CCD array. Independent control is again possible by independently addressing each memory optically. There is also flexibility in a CRIB technique by varying an electric field independently in 3 dimensions, up to the limit imposed by Maxwell’s equations. The ultimate size of a pixel memory is governed by the wavelength. Practical issues will no doubt prevent reaching 1 mode per  $\lambda^2$ , but even 1 mode per  $100\mu\text{m}^2$  would give  $10^4$  modes in a single crystal with input face  $1\text{cm} \times 1\text{cm}$ . The potential for spatial multiplexing appears relatively unexplored in rare-earths, the focus instead being on TBP. It is worth noting that the potential for spatial multiplexing in a solid may be significantly greater than in systems such as gases. This is because in solids there is no need to use a wide beam to avoid effects of lateral diffusion.

This naturally raises the question of compatibility with quantum light sources. There is strong promise of realising these with similar techniques and materials. These techniques will be multi-mode in the same manner (section 6.3).

Creating many memories as described makes much better use of available absorption, but remains somewhat wasteful. Even if we pack memories as closely together as possible, both in frequency and in space, we will still be wasting far more absorbers than we use. Perhaps more importantly, this will impose a significant requirement of time and laser power to maintain the hole-burned spectral structure, particularly in the frequency domain case. The following techniques are less wasteful in the sense of absorption, we note that hole-burning retains significant potential for performance. It also has a unique advantage due to its flexibility. This not only creates a good test-bed for new techniques, but offers a completely different class of device which may be re-programmed to match new needs dynamically.

### 6.2.2 Quiet Optical Rephasing

It is possible to make good use of every single absorber in the inhomogeneous line. This is the case for the original 2PE method of optical rephasing. In a sense this allows control of atoms based not on their position in a crystal or orientation, but directly based on their detuning. The dephasing due to the inhomogeneous broadening can be completely removed.

To avoid the noise on the standard 2PE, the optically-rephased echo must form when the medium is in (or near) the ground state. To achieve this, we rephase the atoms (e.g. with a  $\pi$  pulse), and wait as the atoms reverse their evolution while in the excited state. We must suppress the formation of an echo somehow. We then apply another rephasing pulse, so that now an echo forms in the ground state. By utilising a 3-level scheme, or using phase matched  $\pi$  pulses, it should also be possible to achieve backward retrieval resulting in potentially 100% efficient echoes.

Several options for suppressing echoes have already been mentioned in the context of pulse-sequencing in section 4.4.6. Using a large field to suppress an echo has been demonstrated experimentally by McAuslan et al. [66].

For the more general methods of suppression, it is not completely clear whether spontaneous emission, occurring during the excited state, will leak into the echo mode and prevent its use as a quantum memory.

Further options are to use pulses that are good inversions, but which shuffle the spatial phase of the stored pulse. We have performed proof-of-principle experiments of two variants of this. The first uses  $\pi$  pulses at  $90^\circ$  to the normal propagation direction(6.1). After the first  $\pi$ -pulse, the spatial phase matching is destroyed, while the second  $\pi$ -pulse reinstates it. The second uses pulses that shuffle the phase as a function of frequency, in this case non-optimal chirps.

A limitation of this technique is that a significant delay occurs before the pulse can be retrieved due to the need for rephasing twice. This could be partially avoided using a scheme with shelving on a 3rd level. The 2 rephasing steps are performed first, then the coherence transfer is performed, may be used to make the technique on-demand. The initial waiting period is still required though.

To some extent it should be possible to perform some pulse manipulation optically. Given a long enough time, it is possible to optically manipulate the phase of all ions independently, based on their detuning. The frequency resolution with which this can be done however, is the inverse of the time taken. Thus it should be possible to manipulate the pulse shape by manipulating the phase of the absorbers, but it is unlikely that we can rephase any faster than using a  $\pi$  pulse.

### 6.2.3 Raman Feature with a Spin-Polarised Sample

In order to make use of the entire inhomogeneous line of a crystal, it will be useful, or even necessary, to prepare the entire line within the same ground state, that is to spin-polarise it. When using optical rephasing, it will reduce the number of different Rabi-frequencies such that resonant  $\pi$ -pulses may be used for good rephasing. However a major limitation of using optical rephasing on a large inhomogeneously broadened line is that the ground states are hidden within the line, preventing long-term spin-shelving. Virtual, or Raman transitions may provide a way around this.

Previous demonstrations of such features in rare-earths required complex hole-burning




Figure 6.1: A proof-of-principle, optically rephased echo, expected to meet the noise requirements for quantum memory. The pulse was incident on a feature prepared within a wide spectral trench in PrYSO. The signal shown is of intensity as recovered from a heterodyne measurement. The FID of the (quite narrow) feature can be seen. Two separate  $\pi$  pulses were incident on the sample from the side, i.e. at 90 degrees to the propagation direction of the original pulse. Spikes in the signal (labelled) were registered on the detectors during these as the LO beam was re-routed using an AOM to supply the pulses. This experiment was performed with Sarah Beavan.

steps to create an anti-hole and an empty transition, thus using even less of the available absorption than our present techniques[165]. If the entire line can first be prepared in the one ground state, we may get around this limitation. Given the right parameters of a far off-resonant beam, it should be possible to create narrow features related to the hyperfine structure. Some of these will correspond to absorption and some to gain. Given the spectra-hiding properties of the gradient echo, it should be possible to interact with a single absorption (while hiding the gain features) over an arbitrarily large bandwidth.

The feasibility of spin polarisation reduces rapidly as the inhomogeneous line-width increases, but is aided in crystals (as we expect PrYSO is) where the hole-burning lifetime is limited by spin exchange interactions. This mechanism is largely turned off if all spins become the same. One way it can be achieved naturally is in large fields and low temperatures, then let  $kT$  do the rest, as has been demonstrated with Erbium [106]. Due to its long optical coherence in high fields[166], ErYSO may represent a good near-term application of this technique. Unfortunately the spin-shelving times in this case will only be several 10s of ms due to relaxation back to the lowest ground state.

#### 6.2.4 Hyperfine-Resolved Stoichiometric Crystals

The ideal situation is to have full control over every single absorber. This limit may be theoretically approached by removing inhomogeneous broadening in the crystal. Then, using a technique such as CRIB, ions may be manipulated as desired in frequency space. This would combine the benefits of optical rephasing, where every absorber may contribute, with the flexibility of control when using an external field.

Inhomogeneous line-widths as low as 10MHz were reported by MacFarlane et al.[145, 144] in naturally occurring concentrations ( $\sim 1$ ppm) of  $\text{Nd}^{3+}$  and  $\text{Er}^{3+}$  in  $\text{LiYF}_4$  (YLF). This crystal used isotopically pure Li to reduce the associated broadening. The resultant spectrum consisted of isolated spectral lines corresponding to hyperfine structure and various isotopes of the active ions and their neighbours. Without these impurities, the spectrum should consist of only lines corresponding to the hyperfine transitions.

If these absorption lines are strong enough, such a structure may be used to perform a memory without the normal bandwidth limit due to hyperfine spacing. The key to this is again the spectral 'hiding' which can occur when there is a correlation between inhomogeneous broadening and depth, as in the gradient echo technique (section 4.3.3). A single hyperfine transition (either at the highest or lowest end of the spectrum), if the correlation is applied correctly, will prevent the others from interacting resonantly with the light. There is thus no immediate limit to the broadening that may be applied. If the material hole-burns, it is then a simple matter to prepare it in a single ground state. Coupling to the nuclear spins levels for long term storage is then very easy.

MacFarlane et al. did not report the absorption coefficients for their low inhomogeneously broadened medium. The low concentrations will naturally reduce the optical depth, although this will be compensated to some extent by the reduced line-width. We

expect the optical depth to be reduced though due to the high symmetry of the material used (YLF), such that electric-dipole transitions will remain largely forbidden. The relatively strong magnetic-dipole character of the  $^4I_{15/2} \rightarrow ^4I_{13/2}$  transition in  $\text{Er}^{3+}$  should mean that it's affected less than  $\text{Nd}^{3+}$ , but it will still be affected. In any case, a sufficiently good multi-pass cell or cavity may be sufficient to implement high performance memory.

A prospect with greater potential is to use a material that is stoichiometric in the active ion. In this case the broadening due to doping concentration may be removed completely. Of course the new problem (with an old history of study) is interactions between the active ions. Even if interactions due to proximity may be avoided, such highly absorbing materials will be complicated by optical band-structure[167]. While inter-band transitions should be coherent and potentially usable, they are also mobile. Thus defects or other de-cohering interactions anywhere in a crystal will be sampled rapidly.

The problem is simplified somewhat by using weakly interacting ions in a medium with a large unit cell to maximise their separation. The aim is to reduce de-cohering interactions, and to decrease band structure below the level of hyperfine structure. This is the case for a material such as  $\text{EuCl}_3 \cdot 6\text{H}_2\text{O}$ , which has been re-investigated recently by Ahlefeldt et al. [168]. In this material structure as narrow as 100MHz was observed, with the main source of residual broadening thought to be due to the distribution of chlorine isotopes. This material is expected to have optical depths of around 10<sup>3</sup>dB/mm, around the level of strong interaction at the wavelength level where our current assumptions break down. As the inhomogeneous broadening is decreased, this optical depth will increase. It may be desirable, and also quite simple to temper this in the case of Europium. The two naturally occurring Europium isotopes (151, 153) are spectrally resolvable but chemically identical. Thus by changing the ratio of these isotopes it should be possible to tune the absorption due to the wanted isotope without introducing strain[169].

A final issue that needs addressing is how to induce the required broadening. We would wish to apply broadening on the scale of GHz, which must be reversed rapidly to achieve a gradient echo technique. A large electric field is the experimentally easier choice, but low inhomogeneous broadening and transition dipole moment imply a weak Stark shift. Magnetic fields and Zeeman shifts are an alternative possibility, as are optical fields to induce an AC Stark shift. To obtain large broadening, a pressure shift may be the most practical option. The strength of broadening is relatively easily attained, but the speed of switching is limited by the speed of sound and will limit materials to  $\sim 1\text{mm}$  in order to have  $\sim 1\mu\text{s}$  switching times.

The detuning gradient will probably be necessary to hide unwanted transitions, but the fast switching isn't. Any of the memory techniques we have discussed may be implemented using a detuning-depth correlation to hide structure, such that the detuning field need not be switched at all. As an example, we might create a space-dependent AFC out of an artificially, permanently pressure broadened medium. As long as the echo is retrieved

backwards, the principle of spectral hiding will remain. The convenience and flexibility of rapidly-switched CRIB is desirable however—for the power of mode manipulation, to reduce requirements of auxiliary fields, and to give easy access to ground states.

### 6.3 Memories With and As a Quantum Light Source

Naturally a quantum memory for light is no use without quantum light to go in it. Future REI based memories will need compatibility with sources in other systems, perhaps by wavelength conversion, or to have one implemented using the same system. Wavelength conversion using SPDC has already been demonstrated between Thulium and Erbium wavelengths[2]. To avoid conversion losses and other issues like bandwidth compatibility, we would prefer to avoid the step. Ideally we may have sources based on the same materials and principles as the memory itself.

A source of on-demand quantum light is already essentially a quantum memory. The quantum state might be created directly in the device, rather than input as light, but the readout is the same. Along these lines, there are at least two techniques being investigated for on-demand single photon sources using rare-earths. These are DLCZ[170] and Rephased Amplified Spontaneous Emission (RASE)[171]. The principle of each is very similar—the memory is coerced into a quantum emission, which is measured using a photon counter. The emission event creates a coherence in the same way as an absorption, which may then be rephased. In DLCZ, the memory is stimulated into emitting a photon by an off-resonant pulse, while for RASE the medium put into the excited state so that it emits spontaneously. The RASE process has the advantage of involving rephasing: this means both that its temporally multi-mode, and should generally result in a higher efficiency.

These techniques should work under very similar conditions to quantum memory, ensuring their compatibility in terms of wavelength, bandwidth, and even spatial multiplexing.

The possibilities of combining these techniques with CRIB are also quite interesting, due to the level of control that may be realised. Using either the population inversion (RASE) or Raman scattering (DLCZ), quantum states such as single photons may be created in a CRIB memory. One would simply start with the standard CRIB scheme: either a hole-burned absorption feature with a spatially varying electric field, or a Raman absorption with a spatially varying magnetic field. The detection of spontaneous emission or Raman scattering from the memory then creates the corresponding state inside. This has a particular benefit for DLCZ in that it becomes temporally multi-mode.

Once the state is created in-memory by the detection, it may rephased and manipulated as described in chapter 4. Collapsing the field will suppress the emission, increasing it will increase the bandwidth etc. It should then be possible to manipulate the state's mode, shift it's frequency, suppress events, interfere different events, and perform partial

readouts. This promises a powerful tool to manipulate entanglement.

## 6.4 Summary

Although the main result of this thesis was an efficient quantum memory, an attempt has been made to address the broader issues of ensemble-based quantum memories. We will finish by summarising the problem and the progress that has been made.

For any application it is clear that high efficiency is important, if not crucial. Procedures such as conditional measurements, error-correcting and entanglement distillation make it theoretically possible to perform QIP with lower efficiency, more noisy storage, but the resources needed grow very rapidly. We have demonstrated a 70% efficient memory with a very low noise floor, the first in the right ballpark for the > 90% efficiency estimated as a requirement for repeaters[88]. With simulation we showed how the system should reach over 90% simply by scaling.

Beyond simple high efficiency, the requirement in applications is clear: a combination with high time-bandwidth product (which is also needed to push the efficiency higher), or at least mode-capacity, and long shelving times. The required resources for achieving this in ensembles are also clear: the maximum TBP that can be achieved efficiently is directly proportional to the density of absorbers. To also obtain storage on long time-scales, the system must be controllably coupled to an otherwise isolated shelving state.

The  $4f^N \rightarrow 4f^N$  of REIs in solids, along with their nuclear spins, represent near ideal systems on both counts. This is due to their high shielding and resultant ability to have gas-like spectra while in a solid. The solid state allows much greater absorber densities than otherwise, while maintaining optical access to the long-lived nucleus. Each system is held in a very long-lived natural ion trap. The static nature of the solid state also makes possible large-scale spatial multiplexing as in section 6.2.1. The draw-back of the system is that weak interactions tend to be associated with a weak splitting, strongly limiting the achievable bandwidth. This issue is particularly clear in REI doped solids due to the high resolution possible, but may be expected to show up in other systems. This is simply due to the fact that a weakly split transition requires a lower Q-factor to maintain coherence for a long time.

There will likely be systems where this is less of an issue, and there will clearly be more than one solution, but it appears to be not yet addressed in the literature. If other mechanisms to address the issue are not found, then perhaps the most significant result presented in this thesis is the very simple fact that a detuning gradient may hide resolved spectral structure (section 4.3.3). As discussed, this phenomenon may be used to separate the requirements for high TBP and long storage times. The separation between the main transition and the shelving level will still limit how rapidly the coherence may be shelved, but not the total TBP.

Assuming that memory performing at the right level can be achieved, further barri-



ers to practical applications will be fine control over the stored states, and appropriate quantum light sources. As discussed in chapter 4, a gradient-echo based technique offers the required fine control naturally. The potential for simultaneous partial readout of different components may also prove beneficial to improve entanglement swapping steps in repeaters. The principles of efficiency and manipulation in memories should also prove applicable to quantum light sources using similar techniques.

A final barrier for real-world applicability is not addressed in this work: one of compatibility with other devices. The most promising rare-earth ion, europium, exhibits its useful properties at 580nm. The good points of this wavelength are that high performance light detectors available, and that it's transmitted well in the atmosphere such that satellite communication may be feasible. A major problem in the past has been the availability of good lasers, but this is fading as stabilised solid-state lasers are now coming onto the market. However the fact remains that the standard wavelength for communications is that transmitted in silica fibres.

For compatibility with other systems, and especially to realise repeaters, operation at 1.5 $\mu$ m may be necessary. Erbium remains a very attractive prospect for this reason, even if the potential performance is more limited. It should be possible to realise appropriate performance in terms of TBP, for instance using the method in section 6.2.3. Avenues for long-term shelving also exist (e.g. super-hyperfine splitting or advanced DDC techniques), but the potential to combine them with a high TBP is less clear. Whether it provides quantum storage on its own or not, erbium will almost certainly play a role in conversion.

In the meantime, both Rare-earth ion quantum information and the field of quantum information at large are growing at an ever increasing rate. Whether full-scale applications are realised or not, progress is being made in probing the complex interactions of the solid-state, and the strange limits of quantum mechanics.

---

# Bibliography

---

- [1] J. J. Longdell, E. Fraval, M. J. Sellars, and N. Manson. Stopped Light with Storage Times Greater than One Second Using Electromagnetically Induced Transparency in a Solid. *Physical Review Letters*, 95(6):063601, 2005.
- [2] Erhan Saglamyurek, Neil Sinclair, Jeongwan Jin, Joshua a. Slater, Daniel Oblak, Félix Bussi eres, Mathew George, Raimund Ricken, Wolfgang Sohler, and Wolfgang Tittel. Broadband waveguide quantum memory for entangled photons. *Nature*, 469(7331):512–5, January 2011.
- [3] InternetArchive. Internet Archive: Digital Library of Free Books, Movies, Music and Wayback Machine. (<http://www.archive.org/>).
- [4] Google. Inside Google Books: On the Future of Books (<http://booksearch.blogspot.com/2010/10/on-future-of-books.html>).
- [5] Google. Inside Google Books: Books of the world, stand up and be counted! All 129,864,880 of you. (<http://booksearch.blogspot.com/2010/08/books-of-world-stand-up-and-be-counted.html>).
- [6] Wikimedia. Wikipedia. (<http://www.wikipedia.org>).
- [7] R. MacLeod. *The Library of Alexandria*, volume 65. Tauris, 2005.
- [8] A. M. Turing. On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1):230–265, January 1937.
- [9] Dina Goldin and Peter Wegner. The church-turing thesis: Breaking the myth. *New Computational Paradigms*, pages 152–168, 2005.
- [10] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, June 1982.
- [11] D. Deutsch. Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences (1934-1990)*, 400(1818):97–117, July 1985.
- [12] Richard P. Feynman, Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics, Vol. 3*. Addison Wesley, 1971.

- 
- [13] C.H. Bennett, G. Brassard, and Others. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, volume 175. Bangalore, India, 1984.
- [14] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2010.
- [15] M ukowski, A Zeilinger, and MA Horne. "Event ready Detectors" Bell Experiment for Entanglement Swapping. *Physical Review Letters*, 1993.
- [16] Seth Lloyd and S.L. Braunstein. Quantum Computation over Continuous Variables. *Physical Review Letters*, 82(8):1784–1787, February 1999.
- [17] S.L. Braunstein and Peter van Loock. Quantum information with continuous variables. *Reviews of Modern Physics*, 77(2):513–577, June 2005.
- [18] N. J. Cerf, Gerd Leuchs, and E. S. Polzik. *Quantum information with continuous variables of atoms and light*. Imperial College Press, 2007.
- [19] H J Kimble. The quantum internet. *Nature*, 453(7198):1023–30, June 2008.
- [20] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Reviews of Modern Physics*, 81(3):1301–1350, September 2009.
- [21] AK Ekert. Quantum cryptography based on Bell's theorem. *Physical Review Letters*, 1991.
- [22] 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.
- [23] Zhen-Sheng Yuan, Yu-Ao Chen, Bo Zhao, Shuai Chen, Jörg Schmiedmayer, and Jian-Wei Pan. Experimental demonstration of a BDCZ quantum repeater node. *Nature*, 454(7208):1098–101, August 2008.
- [24] Chin-Wen Chou, Julien Laurat, Hui Deng, Kyung Soo Choi, Hugues de Riedmatten, Daniel Felinto, and H Jeff Kimble. Functional quantum nodes for entanglement distribution over scalable quantum networks. *Science (New York, N.Y.)*, 316(5829):1316–20, June 2007.
- [25] W. Tittel, M. Afzelius, T. Chaneliere, R. Cone, S. Kröll, S.a. Moiseev, and M. J. Sellars. Photon-echo quantum memory in solid state systems. *Laser & Photonics Reviews*, 24(September 2008), 2009.
- [26] Daniel Collins, Nicolas Gisin, and Hugues De Riedmatten \*. Quantum relays for long distance quantum cryptography. *Journal of Modern Optics*, 52(5):735–753, March 2005.

- 
- [27] N. Sangouard, Christoph Simon, Hugues De Riedmatten, and Nicolas Gisin. Quantum repeaters based on atomic ensembles and linear optics. *arXiv*, 0906.2699v, 2009.
- [28] Hans-A. Bachor. *A guide to experiments in quantum optics*. Wiley-VCH, 2nd edition, 1998.
- [29] D.F. Walls and G.J. Milburn. *Quantum Optics*. Springer-Verlag, Berlin, 1994.
- [30] UM Titulaer and RJ Glauber. Density operators for coherent fields. *Physical Review*, 145(4):1041–1050, 1966.
- [31] Brian J Smith and M. G. Raymer. Photon wave functions, wave-packet quantization of light, and coherence theory. *New Journal of Physics*, 9(11):414–414, November 2007.
- [32] Jun John Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. Addison-Wesley, 2010.
- [33] David Jeffery Griffiths. *Introduction to quantum mechanics*. Pearson Prentice Hall, 2005.
- [34] Marlan O. Scully and M. Suhail Zubairy. *Quantum Optics*. Cambridge University Press, 1997.
- [35] A. I. Lvovsky and M. G. Raymer. Continuous-variable optical quantum-state tomography. *Reviews of Modern Physics*, 81(1):299–332, 2009.
- [36] H Hansen, T Aichele, C Hettich, P Lodahl, A. I. Lvovsky, J Mlynek, and S Schiller. Ultrasensitive pulsed, balanced homodyne detector: application to time-domain quantum measurements. *Optics letters*, 26(21):1714–6, November 2001.
- [37] Z. Hradil. Quantum-state estimation. *Physical Review A*, 55(3):1561–1564, 1997.
- [38] Dieter Suter. *The Physics of Atom-Laser Interactions*. Cambridge University Press, 1997.
- [39] David Jeffery Griffiths. *Introduction to electrodynamics*. Prentice Hall, 1999.
- [40] J. Dalibard, J. Dupont-Roc, and C. Cohen-Tannoudji. Vacuum fluctuations and radiation reaction: identification of their respective contributions. *Atoms in electromagnetic fields*, 3:372, 2004.
- [41] Bertrand Delamotte. A hint of renormalization. *American Journal of Physics*, 72(2):170, 2004.
- [42] James M Supplee. Optical Bloch equations: Informal motivation without the Schroedinger equation. *American Journal of Physics*, 68(2):180–185, 2000.

- 
- [43] Leslie Allen and J. H. Eberly. *Optical resonance and two-level atoms*. Courier Dover Publications, 1987.
- [44] F. Bloch, WW Hansen, and M. Packard. Nuclear induction. *Physical review*, 70(7-8):460–474, 1946.
- [45] Colin Warwick. Understanding the Kramers-Kronig Relation Using A Pictorial Proof. (<http://cp.literature.agilent.com/litweb/pdf/5990-5266EN.pdf>).
- [46] G. Compagno, R. Passante, and F. Persico. *Atom-Field Interactions and Dressed Atoms (Cambridge Studies in Modern Optics)*. Cambridge University Press, 1995.
- [47] ID Abella, NA Kurnit, and SR Hartmann. Photon echoes. *Physical Review*, 1966.
- [48] Āaslav Brukner and Anton Zeilinger. Conceptual inadequacy of the Shannon information in quantum measurements. *Physical Review A*, 63(2):1–10, January 2001.
- [49] Robert Griffiths. Types of quantum information. *Physical Review A*, 76(6):1–11, December 2007.
- [50] E. Arthurs and J. Kelly. On the simultaneous measurement of a pair of conjugate observables. *Bell Syst. Tech. J.*, 44:725, 1965.
- [51] Frédéric Grosshans and Philippe Grangier. Quantum cloning and teleportation criteria for continuous quantum variables. *Physical Review A*, 64(1):1–4, 2001.
- [52] WK Wootters and WH Zurek. A single quantum cannot be cloned. *Nature*, 299:802, 1982.
- [53] C.M. Caves. Quantum limits on noise in linear amplifiers. *Physical Review D*, 26(8):1817, 1982.
- [54] S.L. Braunstein and H J Kimble. Teleportation of Continuous Quantum Variables. *Physical Review Letters*, 80(4):869, 1998.
- [55] Vincent Josse, Metin Sabuncu, Nicolas J Cerf, Gerd Leuchs, and Ulrik L. Andersen. Universal Optical Amplification without Nonlinearity. *Physical Review Letters*, 96(163602):1–4, 2006.
- [56] S.L. Braunstein, N.J. Cerf, Sofyan Iblidir, P. Van Loock, and Serge Massar. Optimal cloning of coherent states with a linear amplifier and beam splitters. *arXiv*, 86(21):4938–4941, 2001.
- [57] N. Sangouard, Christoph Simon, M. Afzelius, and Nicolas Gisin. Analysis of a quantum memory for photons based on controlled reversible inhomogeneous broadening. *Physical Review A*, 75(3):032327, 2007.

- 
- [58] J. J. Longdell, G. Hétet, P. 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):1–6, 2008.
- [59] Ryan Camacho, Michael Pack, and John Howell. Slow light with large fractional delays by spectral hole-burning in rubidium vapor. *Physical Review A*, 74(3):1–4, September 2006.
- [60] K.J. Boller, A. Imamolu, and SE Harris. Observation of electromagnetically induced transparency. *Physical review letters*, 66(20):2593, 1991.
- [61] Irina Novikova, Nathaniel Phillips, and Alexey Gorshkov. Optimal light storage with full pulse-shape control. *Physical Review A*, 78(2):021802, 2008.
- [62] M.S. Bigelow, N.N. Lepeshkin, and R.W. Boyd. Observation of ultraslow light propagation in a ruby crystal at room temperature. *Physical Review Letters*, 90(11):113903, 2003.
- [63] NA Kurnit, ID Abella, and SR Hartmann. Observation of a photon echo. *Physical Review Letters*, 13(19):567–568, 1964.
- [64] E.L. Hahn. Spin echoes. *Physical Review*, 80(4):580, 1950.
- [65] DA Wiersma and K. Duppen. Picosecond holographic-grating spectroscopy. *Science*, 1161(1985), 1987.
- [66] 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):1–7, August 2011.
- [67] A. L. Alexander, J. J. Longdell, M. J. Sellars, and N. B. Manson. Photon Echoes Produced by Switching Electric Fields. *Physical Review Letters*, 96(4):043602, February 2006.
- [68] Bjorn Lauritzen, Jiri Minar, Hugues de Riedmatten, M. Afzelius, N. Sangouard, Christoph Simon, and Nicolas Gisin. Telecommunication-Wavelength Solid-State Memory at the Single Photon Level. *Physical Review Letters*, 104(8):080502, 2010.
- [69] A. L. Alexander. *Investigation of Quantum Information Storage in Rare Earth Doped Materials (PhD Thesis)*. PhD thesis, Australian National University, 2007.
- [70] 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(20):2323–5, October 2008.
- [71] Mahdi Hosseini, Ben M. Sparkes, Gabriel Hetet, J. J. Longdell, Ping Koy Lam, and B. C. Buchler. Coherent optical pulse sequencer for quantum applications. *Nature*, 461(7261):241–5, September 2009.

- 
- [72] P. Grangier, J.A. Levenson, and J.P. Poizat. Quantum non-demolition measurements in optics. *Nature*, 396(6711):537–542, September 1998.
- [73] Philippe Grangier, J.A. Levenson, and J.P. Poizat. Quantum non-demolition measurements in optics. *Nature*, 396(6711):537–542, 1998.
- [74] L M Duan, M. D. Lukin, J I Cirac, and P Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature-Suppl.*, 414(6862):413–8, November 2001.
- [75] Haruka Tanji, Saikat Ghosh, Jonathan Simon, Benjamin Bloom, and Vladan Vuletić. Heralded Single-Magnon Quantum Memory for Photon Polarization States. *Physical Review Letters*, 103(4):1–4, 2009.
- [76] Mikael Afzelius and Christoph Simon. Impedance-matched cavity quantum memory. *Physical Review A*, 82(2):1–4, August 2010.
- [77] Thomas Coudreau, Frédéric Grosshans, Samuel Guibal, and Luca Guidoni. Feasibility of a quantum memory for continuous variables based on trapped ions: from generic criteria to practical implementation. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 40(2):413–426, January 2007.
- [78] M. P. Hedges, J. J. Longdell, Yongmin Li, and M. J. Sellars. Efficient quantum memory for light: Supplementary Information. *Nature*, 465(7301):1052–6, June 2010.
- [79] T. Ralph, S. Bartlett, J. O’Brien, G. Pryde, and H. Wiseman. Quantum non-demolition measurements for quantum information. *Physical Review A*, 73(1):012113, 2006.
- [80] R Jozsa. Fidelity for mixed quantum states. *Journal of Modern Optics*, 1994.
- [81] K. Hammerer, M. Wolf, E. Polzik, and J. Cirac. Quantum Benchmark for Storage and Transmission of Coherent States. *Physical Review Letters*, 94(15), 2005.
- [82] N. Cerf, A. Ipe, and X. Rottenberg. Cloning of Continuous Quantum Variables. *Physical Review Letters*, 85(8):1754–1757, August 2000.
- [83] Brian Julsgaard, Jacob Sherson, J. Ignacio Cirac, Jaromir Flurasek, Eugene S. Polzik, and Jaromír Fiurásek. Experimental demonstration of quantum memory for light. *Nature*, 432(7016):482–6, November 2004.
- [84] N. Gisin and S. Massar. Optimal Quantum Cloning Machines. *Physical Review Letters*, 79(11):2153–2156, September 1997.
- [85] K Vogel and H Risken. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Physical Review A*, 1989.

- 
- [86] M. Lobino, C Kupchak, E Figueroa, and A. I. Lvovsky. Memory for light as a quantum process. *Physical review letters*, 102(20):203601, 2009.
- [87] Saleh Rahimi-Keshari, Artur Scherer, Ady Mann, a T Rezakhani, A. I. Lvovsky, and Barry C Sanders. Quantum process tomography with coherent states. *New Journal of Physics*, 13(1):013006, January 2011.
- [88] N. Sangouard, Christoph Simon, Hugues De Riedmatten, Nicolas Gisin, and Hugues de Riedmatten. Quantum repeaters based on atomic ensembles and linear optics. *Review of Modern Physics*, 0906.2699v(1):09062699, March 2009.
- [89] O. Collins, S. Jenkins, a. Kuzmich, and T. Kennedy. Multiplexed Memory-Insensitive Quantum Repeater. *Physical Review Letters*, 98(6):98–101, February 2007.
- [90] Christoph Simon, Hugues de Riedmatten, M. Afzelius, N. Sangouard, Hugo Zbinden, and Nicolas Gisin. Quantum Repeater with Photon Pair Sources and Multimode Memories. *Physical Review Letters*, 98(19):1–4, 2007.
- [91] D R McCamey, J Van Tol, G W Morley, and C Boehme. Electronic spin storage in an electrically readable nuclear spin memory with a lifetime  $\approx 100$  seconds. *Science (New York, N.Y.)*, 330(6011):1652–6, December 2010.
- [92] C. Langer, R. Ozeri, J. Jost, J. Chiaverini, B. DeMarco, a. Ben-Kish, R. Blakestad, J. Britton, D. Hume, W. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaeetz, P. Schmidt, and D. Wineland. Long-Lived Qubit Memory Using Atomic Ions. *Physical Review Letters*, 95(6):2–5, August 2005.
- [93] Rui 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, 2009.
- [94] 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.
- [95] J Becquerel and H.K. Onnes. Letter. *K. Ned. Akad. Wet.*, 10:592, 1908.
- [96] JC Krupa. Optical excitations in lanthanide and actinide compounds. *Journal of Alloys and Compounds*, 225(1-2):1–10, 1995.
- [97] Brian G. Wybourne. *Spectroscopic Properties of Rare Earths*. Interscience Publishers, 1965.
- [98] S. Huefner. *Optical Spectra of Transparent Rare Earth Compounds*. Academic Press, Inc., New York, 1978.



- 
- [99] A. Kaplianskii and R. M. Macfarlane. *Spectroscopy of solids containing rare earth ions*. North-Holland, 1987.
- [100] G. Liu. Electronic Energy Level Structure. In G Liu and B Jacquier, editors, *Spectroscopic Properties of Rare Earths in Optical Materials*, chapter 1. Springer, Berlin, 2005.
- [101] BE Kane. A silicon-based nuclear spin quantum computer. *Nature*, pages 27–40, 1998.
- [102] G.H. Dieke. *Spectra and Energy Levels of Rare Earth Ions in Crystals*. H. Crosswhite, 1968.
- [103] P.S. Peijzel. *Probing high energy levels of lanthanide ions-experiment and theory*. Phd, Universiteit Utrecht, 2004.
- [104] R. M. Macfarlane and R M Shelby. Coherent Transient and Holeburning Spectroscopy. In A. Kaplianskii and R. M. Macfarlane, editors, *Spectroscopy of Solids Containing Rare Earth Ions*, chapter 2. Elsevier, Amsterdam, 1987.
- [105] R. Orbach. Spin-Lattice Relaxation in Rare-Earth Salts. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences (1934-1990)*, 264(1319):458–484, December 1961.
- [106] IN Kurkin and KP Chernov. EPR and spin-lattice relaxation of rare-earth activated centres in Y<sub>2</sub>SiO<sub>5</sub> single crystals. *Physica B+ C*, 101(2):233–238, 1980.
- [107] M. Afzelius, Matthias U. Staudt, Hugues de Riedmatten, Nicolas Gisin, Olivier Guillot-Noël, Philippe Goldner, Robert Marino, Pierre Porcher, Enrico Cavalli, and Marco Bettinelli. Efficient optical pumping of Zeeman spin levels in Nd<sup>3+</sup>:YVO<sub>4</sub>/Nd<sup>3+</sup>:YVO<sub>4</sub>. *Journal of Luminescence*, 130(9):1566–1571, September 2010.
- [108] Nicklas Ohlsson, Mattias Nilsson, Stefan Kröll, and R Krishna Mohan. Long-time-storage mechanism for Tm:YAG in a magnetic field. *Optics letters*, 28(6):450–2, March 2003.
- [109] R. M. Macfarlane. High-resolution laser spectroscopy of rare-earth doped insulators: a personal perspective. *Journal of Luminescence*, 100(1-4):1–20, December 2002.
- [110] Flurin Könz, Y. Sun, C. Thiel, R. Cone, R. Equall, R. Hutcheson, and R. M. Macfarlane. Temperature and concentration dependence of optical dephasing, spectral-hole lifetime, and anisotropic absorption in Eu<sup>3+</sup>:Y<sub>2</sub>SiO<sub>5</sub>. *Physical Review B*, 68(8):1–9, August 2003.
- [111] B.R. Judd. Optical absorption intensities of rare-earth ions. *Physical Review*, 127(3):750, 1962.

- 
- [112] GS Ofelt. Intensities of Crystal Spectra of Rare-Earth Ions. *The Journal of Chemical Physics*, 37(3):511, 1962.
- [113] B Wybourne. The fascination of the rare earths. Then, now and in the future. *Journal of Alloys and Compounds*, 380(1-2):96–100, October 2004.
- [114] John M. Hayes, Roland P. Stout, and Gerald J. Small. Hole burning, soft-phonon scattering and anomalously fast dephasing of impurity electronic transitions in glasses. *The Journal of Chemical Physics*, 73(8):4129, 1980.
- [115] R. Van den Berg and S. Volker. Optical hole-burning in organic glasses. *Journal of Luminescence*, 38(1-6):25–30, December 1987.
- [116] W.B. Mims. Electron Spin Echoes. In Stanley Geschwind, editor, *Electron paramagnetic resonance*, page 263. Plenum Press, 1972.
- [117] RM Shelby and CS Yannoni. Optically detected coherent transients in nuclear hyperfine levels. *Physical Review Letters*, 41(25):1739–1742, 1978.
- [118] J. J. Longdell, M. J. Sellars, and N. Manson. Demonstration of Conditional Quantum Phase Shift Between Ions in a Solid. *Physical Review Letters*, 93(13):1–4, September 2004.
- [119] Lorenza Viola and Seth Lloyd. Dynamical suppression of decoherence in two-state quantum systems. *Physical Review A*, 58(4):2733–2744, October 1998.
- [120] SC Rand, A Wokaun, RG DeVoe, and RG Brewer. Magic-angle line narrowing in optical spectroscopy. *Physical Review Letters*, 43(25):1868–1871, 1979.
- [121] R. M. Macfarlane, CS Yannoni, and RM Shelby. Optical line narrowing by nuclear spin decoupling in Pr<sup>3+</sup>: LaF<sub>3</sub>. *Optics Communications*, 32(1):101–104, 1980.
- [122] 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, 2005.
- [123] J. J. Longdell. *Quantum Information Processing in Rare Earth Ion Doped Insulators*. Phd, Australian National University, 2003.
- [124] G. Pryde. *Ultra-high resolution spectroscopic studies of optical dephasing in solids*. Phd, Australian National University, 1999.
- [125] Y. Sun. Rare Earth Materials in Optical Storage and Data Processing Applications. In G. Liu and B Jacquier, editors, *Spectroscopic Properties of Rare Earths in Optical Materials*, chapter 7. Springer, Berlin, 2005.
- [126] C.W. Thiel, Thomas Böttger, and R.L. Cone. Rare-earth-doped materials for applications in quantum information storage and signal processing. *Journal of Luminescence*, 131(3):353–361, March 2011.

- 
- [127] Thomas Böttger, C. Thiel, R. Cone, and Y. Sun. Effects of magnetic field orientation on optical decoherence in  $\text{Er}^{3+}:\text{Y}_2\text{SiO}_5$ . *Physical Review B*, 79(11):115104, 2009.
- [128] S. Hastings-Simon, B. Lauritzen, M. Staudt, J. van Mechelen, C. Simon, H. de Riedmatten, M. Afzelius, and N. Gisin. Zeeman-level lifetimes in  $\text{Er}^{3+}:\text{Y}_2\text{SiO}_5$ . *Physical Review B*, 78(8):1–7, August 2008.
- [129] R. M. Macfarlane. Spectral hole burning in the trivalent thulium ion. *Optics letters*, 18(10):829–31, May 1993.
- [130] R. Equall, Y. Sun, R. Cone, and R. M. Macfarlane. Ultraslow optical dephasing in  $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$ . *Physical Review Letters*, 72(14):2179, 1994.
- [131] J. P. Elliott, B. R. Judd, and W. A. Runciman. Energy Levels in Rare-Earth Ions. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 240(1223):509–523, July 1957.
- [132] Ryuzi Yano, Masaharu Mitsunaga, and Naoshi Uesugi. Ultralong optical dephasing time in  $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$ . *Optics letters*, 16(23):1884, 1991.
- [133] A. L. Alexander, J. J. Longdell, and M. J. Sellars. Measurement of the ground-state hyperfine coherence time of  $\text{EuYSO}$ . *Journal of the Optical Society of America B*, 24(9):2479, 2007.
- [134] J. J. Longdell, A. L. Alexander, and M. J. Sellars. Characterization of the hyperfine interaction in europium-doped yttrium orthosilicate and europium chloride hexahydrate. *Physical Review B*, 74(19):195101, 2006.
- [135] Felix R. Graf. *Investigation of spectral dynamics in rare earth ion doped crystals using high resolution laser techniques*. PhD thesis, Swiss Federal Institute of Technology, Zurich, 1998.
- [136] G. Hétet, J. J. Longdell, A. L. Alexander, P K Lam, and M. J. Sellars. Electro-optic quantum memory for light using two-level atoms. *Physical review letters*, 100(2):023601, January 2008.
- [137] Mattias Nilsson, Lars Rippe, and Stefan Kröll. Hole-burning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids demonstrated in  $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ . *Physical Review B*, 70(21):1–12, December 2004.
- [138] J. J. Longdell, M. J. Sellars, and N. B. Manson. Hyperfine interaction in ground and excited states of praseodymium-doped yttrium orthosilicate. *Physical Review B*, 66(3):035101, June 2002.
- [139] M. Nilsson. *Coherent Interactions in Rare-Earth-Ion-Doped Crystals for Applications in Quantum Information Science*. Phd, Lund, 2004.

- 
- [140] S. Beavan, E. Fraval, M. J. Sellars, and J. J. Longdell. Demonstration of the reduction of decoherent errors in a solid-state qubit using dynamic decoupling techniques. *Physical Review A*, 80(3):1–9, September 2009.
- [141] Felix R Graf, Alois Renn, Urs P Wild, and Masaharu Mitsunaga. Site interference in Stark-modulated photon echoes. *Physical Review B*, 55(17):225–229, 1997.
- [142] A. Winnacker, R.M. Shelby, and R. M. Macfarlane. Photon-gated hole burning: a new mechanism using two-step photoionization. *Optics letters*, 10(7):350–352, 1985.
- [143] R. M. Macfarlane. Photon-gated spectral holeburning. *Journal of Luminescence*, 38(1-6):20–24, December 1987.
- [144] R. M. Macfarlane, R. Meltzer, and B. Malkin. Optical measurement of the isotope shifts and hyperfine and superhyperfine interactions of Nd in the solid state. *Physical Review B*, 58(9):5692–5700, September 1998.
- [145] R. M. Macfarlane, A. Cassanho, and RS Meltzer. Inhomogeneous broadening by nuclear spin fields: a new limit for optical transitions in solids. *Physical review letters*, 69(3):542–545, 1992.
- [146] V Damon, M Bonarota, A. Louchet-Chauvet, T. ChaneliÅ\cLre, and J.L.L. GouÅt. Revival of Silenced Echo and Quantum Memory for Light. *Arxiv preprint arXiv:1104.4875*, 2011.
- [147] 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, 2008.
- [148] 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(20):2323–5, October 2008.
- [149] 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):023803, 2008.
- [150] G. Collett, P. D. Drummond, P. Cochrane, G. Dennis, and J. J. Hope. eXtensible Multi-Dimensional Simulator.
- [151] 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):023803, 2008.
- [152] M. D. Levenson. *Introduction to Nonlinear Laser Spectroscopy*. Academic Press, Inc., New York, 1982.

- 
- [153] 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.
- [154] M Mitsunaga. Time-domain optical data storage by photon echo. *Optical and quantum electronics*, 24:1137–1150, 1992.
- [155] A Delfan Abazari, E Saglamyurek, and R Ricken. State transformation in photon-echo quantum memory. *arXiv*, 2009.
- [156] M. J. Sellars. *Ultra-High Resolution Laser Spectroscopy of Rare Earth Doped Solids*. Phd, Australian National University, 1995.
- [157] R. Orbach. Spin-Lattice Relaxation in Rare-Earth Salts: Field Dependence of the Two-Phonon Process. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences (1934-1990)*, 264(1319):485–495, December 1961.
- [158] M. J. Sellars, G.J Pryde, N.B Manson, and E.R Krausz. 50kHz absorption line in Y<sub>2</sub>SiO<sub>5</sub>:Eu<sup>3+</sup>. *Journal of Luminescence*, 87-89:833–835, May 2000.
- [159] M. J. Sellars, R. S. Meltzer, P. T. H. Fisk, and N. B. Manson. Time-resolved ultranarrow optical hole burning of a crystalline solid: Y<sub>2</sub>O<sub>3</sub>:Eu<sup>3+</sup>. *Journal of the Optical Society of America B*, 11(8):1468, August 1994.
- [160] D. T. Smithey, M. Beck, and M. G. Raymer. Measurement of the Wigner Distribution and Density Matrix of a Light Mode Using Optical Homodyne Tomography: Application to Squeezed States and the Vacuum. *Phys. Rev. Lett*, 70(9), 1993.
- [161] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- [162] B Ham, P Hemmer, and M Shahriar. Efficient electromagnetically induced transparency in a rare-earth doped crystal. *Optics Communications*, 144(4-6):227–230, December 1997.
- [163] Elliot Fraval. *Minimising the Decoherence of Rare Earth Ion Solid State Spin Qubits*. Phd, Australian National University, 2006.
- [164] Michael Varnava, Daniel Browne, and Terry Rudolph. Loss Tolerance in One-Way Quantum Computation via Counterfactual Error Correction. *Physical Review Letters*, 97(12):120501, 2006.
- [165] a. Turukhin, V. Sudarshanam, M. Shahriar, J. Musser, B. Ham, and P. Hemmer. Observation of Ultraslow and Stored Light Pulses in a Solid. *Physical Review Letters*, 88(2):1–4, December 2001.

- 
- [166] Y. Sun. Recent progress in developing new rare earth materials for hole burning and coherent transient applications. *Journal of Luminescence*, 98(1-4):281–287, July 2002.
- [167] R. L. Cone and R. S. Meltzer. Ion-Ion Interactions and Exciton Effects in Rare Earth Insulators. In A. Kaplanski and R. M. Macfarlane, editors, *Spectroscopy of Solids Containing Rare Earth Ions*, chapter 8. Elsevier, Amsterdam, 1987.
- [168] R L Ahlefeldt, A Smith, and M. J. Sellars. Ligand isotope structure of the optical  $7F_0/5D_0$  transition in  $\text{EuCl}_3 \cdot 6\text{H}_2\text{O}$ . *Physical Review B*, 80(20):205106, 2009.
- [169] M. J. Sellars. Private Communication.
- [170] Elizabeth Goldschmidt, Sarah Beavan, Matthew Eisaman, Jingyun Fan, Michael Hohensee, Zachary Levine, Ludwig Mathey, Sergey Polyakov, and Alan Migdall. State Preparation for Rare-Earth Ion Based Quantum Memory. *APS Division of Atomic, Molecular and Optical Physics Meeting Abstracts*, -1:1001, May 2009.
- [171] Patrick M. Ledingham, William R. Naylor, and J. J. Longdell. Nonclassical photon streams using rephased amplified spontaneous emission. *Physical Review A*, 81(1), January 2010.